arXiv:1412.0367v1 [stat.AP] 1 Dec 2014

Bayesian nonparametric modeling for mean residual life regression

Valerie Poynor and Athanasios Kottas *

Abstract: The mean residual life function is a key functional for a survival distribution. It has practically useful interpretation as the expected remaining lifetime given survival up to a particular time point, and it also characterizes the survival distribution. However, it has received limited attention in terms of inference methods under a probabilistic modeling framework. In this paper, we seek to provide general inference methodology for mean residual life regression. Survival data often include a set of predictor variables for the survival response distribution, and in many cases it is natural to include the covariates as random variables into the modeling. We thus propose a Dirichlet process mixture modeling approach for the joint stochastic mechanism of the covariates and survival responses. This approach implies a flexible model structure for the mean residual life of the conditional response distribution, allowing general shapes for mean residual life as a function of covariates given a specific time point, as well as a function of time given particular values of the covariate vector. To expand the scope of the modeling framework, we extend the mixture model to incorporate dependence across experimental groups, such as treatment and control groups. This extension is built from a dependent Dirichlet process prior for the group-specific mixing distributions, with common locations and weights that vary across groups through latent bivariate beta distributed random variables. We develop properties of the proposed regression models, and discuss methods for prior specification and posterior inference. The different components of the methodology are illustrated with simulated data sets. Moreover, the modeling approach is applied to a data set comprising right censored survival times of patients with small cell lung cancer.

KEY WORDS: Dirichlet process mixture models; Dependent Dirichlet process; Markov chain Monte Carlo; Mean residual life; Survival regression analysis.

*V. Poynor is Postdoctoral Researcher, Southwest Fisheries Science Center, NOAA, Santa Cruz, CA 95060 (E-mail: *vpoynor@soe.ucsc.edu*), and A. Kottas is Professor of Statistics, Department of Applied Mathematics and Statistics, University of California, Santa Cruz, CA 95064. (E-mail: *thanos@ams.ucsc.edu*). This research is part of the Ph.D. dissertation of Valerie Poynor, completed at University of California, Santa Cruz, and was supported in part by the National Science Foundation under award DMS 1310438.

1 Introduction

The mean residual life (mrl) function of a continuous random variable T, with positive values and finite mean, μ, provides the expected remaining lifetime given survival up to time t:

$$m(t) \;=\; \frac{\int_t^\infty S(u)\mathrm{d}u}{S(t)} = \frac{\mu - \int_0^t S(u)\mathrm{d}u}{S(t)} \tag{1}$$

where the latter expression is achieved using $\mu = \int_0^\infty S(u)\mathrm{d}u$. From (1), it is clear that the mrl function is defined through the survival function, $S(t) = \Pr(T > t)$, however, the survival function can also be defined through the mrl function via the Inversion formula: $S(t) = m(0)/m(t)\exp\left[-\int_0^t 1/m(u)\mathrm{d}u\right]$ (Hall and Wellner, 1981). Thus the mrl function has an one-to-one relationship with the survival distribution. In addition to its unique theoretical properties, the mrl function is of practical importance in a variety of fields, such as reliability, medicine, and actuarial science. For example, in reliability analysis of interest is estimation of the time to failure of a machine, given the machine has not failed up to a particular time. In medicine, doctors wish to study the life expectancy of a patient who has been battling cancer for a period of time. Often associated with the survival times in the aforementioned scenarios, there are a set of covariates (e.g., age, gender, mass, etc.). In the regression setting, it is of interest to develop modeling that allows flexible mrl function shapes over the covariate space. To our knowledge, there is no work in the Bayesian nonparametric literature that explores modeling and inference for the mrl function in the presence of covariates. Here, we propose general inference methodology for mrl regression built from nonparametric mixture modeling for the joint stochastic mechanism of covariates and survival responses.

Modeling approaches for Bayesian survival regression can be broadly categorized into fully parametric and semiparametric; see, e.g., Ibrahim et al. (2001). We do not attempt a full review here, but note that parametric approaches typically build from a parametric survival distribution, such as the exponential, Weibull, gamma, or lognormal, using a link function on one of the parameters to equate a linear structure on the set of covariates, $\mathbf{x}$, and their effects, $\boldsymbol{\beta}$. Regarding semiparametric survival regression models, the main emphasis of the Bayesian literature has been on the proportional hazards (PH) and accelerated failure time (AFT) specifications. Under the former, the hazard function of the survival responses is factorized into a covariate-free baseline hazard, $h_0(t)$, and a non-negative function of the covariates, $c(\mathbf{x}'\boldsymbol{\beta})$. Equivalently, the PH model postulates the form $S_0(t)^{c(\mathbf{x}'\boldsymbol{\beta})}$ for the response survival function, where $S_0(t)$ is the baseline survival function. The AFT specification utilizes the traditional linear regression structure for the survival responses on the logarithmic scale, $\log(t) = \mathbf{x}'\boldsymbol{\beta} + \epsilon$. Hence, under

the AFT setting, the response survival function is expressed as $S_0(t\exp(-\mathbf{x}'\boldsymbol{\beta}))$, where $S_0(t)$ is the survival function of the errors on the original scale. The Bayesian nonparametrics literature includes several methods for semiparametric PH and AFT modeling based on different types of nonparametric priors for h_0 or S_0, respectively, and parametric priors for $\boldsymbol{\beta}$. A sample of related references is Gelfand and Mallick (1995), Ibrahim et al. (1999), and Hanson et al. (2012) under the PH setting, and Walker and Mallick (1999), Kottas and Gelfand (2001), Hanson and Johnson (2002), and Argiento et al. (2009) for AFT semiparametric modeling.

Semiparametric specifications for survival regression provide for ready interpretation of the model components and are therefore attractive to practitioners. However, the structured model scenarios they postulate are not supported by several data sets, and this is well documented in the literature for both the PH and AFT settings. In the context of Bayesian modeling, the work by DeIorio et al. (2009) is the only attempt we are aware of to relax semiparametric specifications while retaining a nonparametric, fully inferential framework. DeIorio et al. (2009), extending earlier work in DeIorio et al. (2004), develop a mixture model for survival responses (on a logarithmic scale) with a covariate-dependent mixing distribution modeled with a linear dependent Dirichlet process prior. A more detailed description of this model is given in Section 4.2, we simply note here that, although it allows for flexible response distributions, it can not capture general, non-linear regression function shapes.

Our objective is to develop a modeling framework that enables flexible inference for both the conditional response survival distribution and for regression relationships, with particular emphasis on mrl regression functions. To this end, we propose nonparametric mixture modeling for the joint stochastic mechanism of the covariates and survival responses, from which inference for mrl regression emerges through the implied conditional response distribution. This curve fitting approach to regression was proposed by Müller et al. (1996) for real-valued responses, and was extended with respect to the resulting inferences by Taddy and Kottas (2010); an overview and additional references are given in Kottas et al. (2013). For problems with a small to moderate number of random covariates, this modeling approach is attractive in terms of its inferential flexibility. Here, we argue for its utility for appropriate applications in the biomedical sciences. At the same time, survival data typically comprise responses (and associated covariates) from subjects assigned to different experimental groups, such as control and treatment groups. The treatment indicator can not be meaningfully incorporated into the joint response-covariate mixture modeling framework as an additional component of the mixture kernel. We thus extend the model to allow distinct mixture distributions for the different groups, which are however dependent in the prior with the dependence built in a nonparametric fashion. We develop this extension in the context of two groups, using a novel dependent Dirichlet process prior for the

group-specific mixing distributions. A key aspect of the modeling approach is the choice of the mixture kernel that corresponds to the survival responses. Moreover, even though we do not model directly the mrl function of the response distribution, we show that the implied model for the mrl function given the covariates has an appealing structure as a locally weighted mixture of the kernel mrl functions, with weights that depend on both time and the covariates.

The outline of the paper is as follows. Section 2 develops the implied conditional regression modeling approach, including the mixture model formulation, approaches to prior selection and posterior inference, and a simulation data example. In Section 3, we present the model elaboration to incorporate survival data from different experimental groups. We study properties of the proposed dependent Dirichlet process prior model, and discuss posterior simulation under the full hierarchical regression model, and two examples with synthetic data. In Section 4, we provide a detailed analysis of a standard data set from the literature on right censored survival times for patients with small lung cancer, including formal comparison with simpler Bayesian models. Section 5 includes a summary, and the two appendixes collect technical details on implementation of posterior inference and computation of the model comparison criterion.

2 Curve Fitting with Random Covariates

When a covariate can be considered to be random, by which we mean the covariate is not fixed such as a patient being assigned to treatment or control groups, it makes sense to model the covariate jointly with the survival response variable. The benefit of this modeling approach is the simplicity of obtaining any conditional or marginal distributions and functionals of interest. We show that functionals such as the mean regression and mrl regression have an interpretable form under the model framework. In addition, we are not restricted to any particular shape in regards to functionals within and across the covariate space.

2.1 Model formulation

Let $\boldsymbol{x}$ be a vector of random covariates and $t > 0$ the survival time of a subject. We model the joint response-covariate density using a Dirichlet process mixture model (DPMM), $f(t, \boldsymbol{x}; G) = \int_{\boldsymbol{\Theta}} k(t, \boldsymbol{x}; \boldsymbol{\theta}) dG(\boldsymbol{\theta})$. Recall that the DP is a stochastic process with random sample paths that are distributions (Ferguson, 1973). The DP is defined through a baseline distribution, G_0 and the precision parameter α. The larger the value of α, the closer the sample path will resemble the baseline distribution. In the DPMM, the DP prior is placed on the mixing distribution $G(\cdot)$. We will work with the truncated version of the SB constructive definition (Sethuraman, 1994) of the DP, $G(\cdot) \approx G_L(\cdot) = \sum_{l=1}^{L} p_l \delta_{\boldsymbol{\theta}_l}(\cdot)$, where $\boldsymbol{\theta}_l \overset{iid}{\sim} G_0$ for $l = 1, ..., L$, and $p_1 = v_1$, $p_l =$

$v_l \prod_{r=1}^{l-1}(1-v_r)$, where $v_r \overset{iid}{\sim} Beta(1,\alpha)$ for $r = 1, ..., L-1$, and $p_L = 1 - \sum_{l=1}^{L-1} p_l$. Thus, the model is given by:

$$f(t, \boldsymbol{x}; G) = \int_{\boldsymbol{\Theta}} k(t, \boldsymbol{x}; \boldsymbol{\theta}) dG(\boldsymbol{\theta}) \approx \sum_{l=1}^{L} p_l k(t, \boldsymbol{x}; \boldsymbol{\theta}_l) \tag{2}$$

Under this model structure, we obtain an analogous mixture formation for the corresponding regression functionals. The mean regression can be expressed as a weighted mixture of the means associated with the conditional kernel components. Specifically, the form of the mean regression is given by:

$$\begin{aligned} E(t|\boldsymbol{x}_0; G_L) &= \frac{\int_0^\infty t f(t, \boldsymbol{x}_0; G_L) dt}{f(\boldsymbol{x}_0; G_L)} = \frac{\int_0^\infty t \int_{\boldsymbol{\Theta}} k(t, \boldsymbol{x}_0; \boldsymbol{\theta}) dG_L(\boldsymbol{\theta}) dt}{f(\boldsymbol{x}_0; G_L)} \\ &= \frac{\int_{\boldsymbol{\Theta}} \int_0^\infty t k(t, \boldsymbol{x}_0; \boldsymbol{\theta}) dt dG_L(\boldsymbol{\theta})}{f(\boldsymbol{x}_0; G_L)} = \frac{\sum_{l=1}^{L} p_l \int_0^\infty t k(t, \boldsymbol{x}_0; \boldsymbol{\theta}_l) dt}{\sum_{l=1}^{L} p_l k(\boldsymbol{x}_0; \boldsymbol{\theta}_l)} \\ &= \sum_{l=1}^{L} q_l(\boldsymbol{x}_0; \boldsymbol{\theta}_l) E(t|\boldsymbol{x}_0; \boldsymbol{\theta}_l) \end{aligned} \tag{3}$$

where $q_l(\boldsymbol{x}_0; \boldsymbol{\theta}_l) = p_l k(\boldsymbol{x}_0; \boldsymbol{\theta}_l) / \{\sum_{l=1}^{L} p_l k(\boldsymbol{x}_0; \boldsymbol{\theta}_l)\}$ are covariate dependent weights (Müller et al., 1996). These covariate dependent weights illustrate the potential for seeing a non-standard relationship of the mean regression across the covariate space. Similarly, the mrl function can be written as,

$$m(t|\boldsymbol{x}_0; G_L) = \sum_{l=1}^{L} q_l(t, \boldsymbol{x}_0; \boldsymbol{\theta}_l) m(t|\boldsymbol{x}_0; \boldsymbol{\theta}_l) \tag{4}$$

where $q_l(\boldsymbol{x}_0; \boldsymbol{\theta}_l) = p_l k(\boldsymbol{x}_0; \boldsymbol{\theta}_l) S(t|\boldsymbol{x}_0; \boldsymbol{\theta}_l) / \{\sum_{l=1}^{L} p_l k(\boldsymbol{x}_0; \boldsymbol{\theta}_l) S(t|\boldsymbol{x}_0; \boldsymbol{\theta}_l)\}$. Therefore, we can think of the mrl regression function as a finite weighted sum of the mrl functions associated with the conditional kernel components, with weights that are dependent on time as well as the covariate values. Aside from the nice interpretation of the form of the mrl regression function under this model, (4) shows the capacity of the model to capture non-standard relationships across the covariate space as well as unique features within the mrl regression function.

As Poynor and Kottas (2014) points out, the mean of the model must exist and be finite in order for the mrl function to exist. They demonstrate the importance of the kernel choice in the DPMM to ensure finiteness of the mean. In the regression setting, we are interested in the conditional mrl at any fixed set of covariates, $m(t|\boldsymbol{x}_0; G_L)$. The sufficiency condition that ensures the finiteness for the mean discussed in Poynor and Kottas (2014), can be extended

for $E(t|\boldsymbol{x}_0; G_L)$. Let $T(\boldsymbol{x}_0, \boldsymbol{\theta}) = \int_0^\infty tk(t; \boldsymbol{x}_0; \boldsymbol{\theta})dt$ where $k(t, \boldsymbol{x}_0; \boldsymbol{\theta}) = k(t|\boldsymbol{x}_0; \boldsymbol{\theta})k(\boldsymbol{x}_0; \boldsymbol{\theta})$. The condition states that if $A(\boldsymbol{x}_0, \boldsymbol{\psi}) = \int_{\boldsymbol{\Theta}} T(\boldsymbol{x}_0, \boldsymbol{\theta})dG_0(\boldsymbol{\theta}) < \infty$ for all $\boldsymbol{\theta}$, then $E(t|\boldsymbol{x}_0; G_L) < \infty$, almost surely. If we chose independent kernel distributions for T and $\boldsymbol{X}_0$, we have $k(t, \boldsymbol{x}_0; \boldsymbol{\theta}) = k(t|\boldsymbol{x}_0; \boldsymbol{\theta})k(\boldsymbol{x}_0; \boldsymbol{\theta}) = k(t; \boldsymbol{\theta})k(\boldsymbol{x}_0; \boldsymbol{\theta})$. We need only choose $k(\boldsymbol{x})$ such that the support is consistent with the support of the covariates. In the case where the covariates are continuous and take values on the real line (possibly after transformation), the multivariate normal distribution for $k(\boldsymbol{x})$ is a natural choice. Turning to $k(t)$, under the independent scenario, $A(\boldsymbol{x}_0, \boldsymbol{\psi})$ becomes the same with $A(\boldsymbol{\psi})$ in Poynor and Kottas (2014). In their discussion, the gamma distribution was the clear winner for the kernel choice, completing justification for the following kernel distribution: $k(t, \boldsymbol{x}_0|\boldsymbol{\theta}) = \Gamma(t|\theta, \phi)N_d(\boldsymbol{x}|\beta, \kappa^2)$, where d is the number of covariates. In our simulation example, we consider a single continuous real-valued covariate ($d = 1$). Let t_i be the survival time and x_i be the corresponding real valued covariate for subject i, for $i = 1, ..., n$. We consider the following DPMM,

$$\begin{aligned} t_i, x_i|\boldsymbol{\theta}, \mathrm{w}_i &\overset{ind}{\sim} \Gamma(t_i; e^{\theta_{\mathrm{w}_i}}, e^{\phi_{\mathrm{w}_i}})N(x_i|\beta_{\mathrm{w}_i}, \kappa^2_{\mathrm{w}_i}) \\ \mathrm{w}_i|\mathbf{p} &\overset{iid}{\sim} \sum_{l=1}^{L} p_l\delta_l(\mathrm{w}_i) \\ (\theta_l, \phi_l, \beta_l, \kappa_l^2)'|\boldsymbol{\mu}, \boldsymbol{\Sigma}, \lambda, \tau^2, \rho &\overset{iid}{\sim} N_2((\theta_l, \phi_l)'|\boldsymbol{\mu}, \boldsymbol{\Sigma})N(\beta_l|\lambda, \tau^2)\Gamma^{-1}(\kappa_l^2|a, \rho) \end{aligned} \tag{5}$$

where $\boldsymbol{p}|\alpha \sim f(\boldsymbol{p}|\alpha) = \alpha^{L-1}p_L^{\alpha-1}(1-p_1)^{-1}(1-(p_1+p_2))^{-1} \times ... \times (1-\sum_{l=1}^{L-2} p_l)^{-1}$ is a special case of the generalized Dirichlet distribution (Conner and Mosemann, 1969). We place the following priors: $\alpha \sim \Gamma(\alpha|a_\alpha, b_\alpha(\text{rate}))$, $\boldsymbol{\mu} \sim N_2(\boldsymbol{\mu}|a_\mu, B_\mu)$, $\boldsymbol{\Sigma} \sim IWish(\boldsymbol{\Sigma}|a_\Sigma, B_\Sigma)$, $\lambda \sim N(\lambda|a_\lambda, b_\lambda)$, $\tau^2 \sim \Gamma^{-1}(\tau^2|a_\tau, b_\tau)$, and $\rho \sim \Gamma(\rho|a_\rho, b_\rho)$. The MCMC is relatively straight forward, only requiring one Metropolis-Hastings step for the parameters of the gamma kernel, just as in the model with no covariates. The rest of the parameters can be sampled via Gibbs steps.

Although, we consider only a single continuous covariate having an independent kernel composition with the survival times, we may also consider discrete covariates and more general kernel structures. Discrete covariates, such as indication of male or female, number of cystic masses, level of pain intensity, are very common in survival analysis. One way to extend the model to include discrete covariates, $\boldsymbol{w}_i$, is to introduce an appropriate density for the kernel product, such that the product kernel becomes $k(t|\boldsymbol{\theta}_T)k(\boldsymbol{x}|\boldsymbol{\theta}_X)k(\boldsymbol{w}|\boldsymbol{\theta}_W)$, where $k(\boldsymbol{w}|\boldsymbol{\theta}_W)$ can be a product of densities such that the categorical covariates are assumed independent in the kernel. In the case where the covariate has a finite upper bound, such as a pain level scale, the binomial is the obvious choice. When the covariate support has no upper bound, the Poisson or negative binomial distribution may be considered.

For added model flexibility, we may seek to incorporate a dependency between the covariates and the survival times within the kernel. A structured approach to build dependence would be to write the kernel as a product of the conditional distribution of the survival times on the covariate and the marginal of the covariates, $k(t|\boldsymbol{x})k(\boldsymbol{x})$. We have made a clear argument for modeling the survival times with a gamma kernel, so one possibility is to incorporate the covariates in one of the parameters of a gamma distribution for $k(t|\boldsymbol{x})$. For example we can write the covariates as a linear combination such as $k(t|\boldsymbol{x}) = \Gamma(t|exp(\theta), exp(\boldsymbol{x}^T\boldsymbol{\beta}))$, such that $E(T|\boldsymbol{x}) = exp(\theta - \boldsymbol{x}^T\boldsymbol{\beta})$. An appropriate marginal kernel, $k(\boldsymbol{x})$, will then complete the joint kernel.

Note that under the curve fitting approach, as the number of covariates increases, the more computationally expensive fitting the model becomes. This is due mainly to the dimension of the covariance matrix for the random covariates in the kernel $\boldsymbol{\Sigma}$. If d is the number of random covariates, the number of parameters in $\boldsymbol{\Sigma}$ that will have to be updated is $Ld(d+1)/2$. One can see how quickly this number can grow with growing d. Thus, curve fitting under the DP mixture framework is best suited for situations in which the number of random covariates is small to moderate.

2.2 Prior selection and posterior inference

For prior specification, we use the same ideas discussed in Poynor and Kottas (2014). In particular, by means of imagining one component covering the prior range believed by the expert, here, for both the survival responses and the covariates. Under the product kernel, we specify the prior parameters associated with the survival times independently of the prior parameters for the covariate values.

We can obtain posterior point and interval estimates for the desired survival functionals by computing the value of the functional at each posterior sample of the parameters over a grid of survival times at a particular value of the covariate, x_0, and saving the desired quantiles. The expressions of the density, survival, hazard, and mrl functions are given respectively below:

$$
\begin{aligned}
f(t|\boldsymbol{x}_0; G_L) &= \frac{f(t, \boldsymbol{x}_0; G_L)}{f(\boldsymbol{x}_0; G_L)} = \frac{\sum_{l=1}^{L} p_l k(t, \boldsymbol{x}_0; \boldsymbol{\theta}_l)}{\sum_{l=1}^{L} p_l k(\boldsymbol{x}_0; \boldsymbol{\theta}_l)} \\
S(t|\boldsymbol{x}_0; G_L) &= 1 - \int_0^t f(u|\boldsymbol{x}_0, G_L)du = 1 - \frac{\sum_{l=1}^{L} p_l k(\boldsymbol{x}_0; \boldsymbol{\theta}_l) K(t|\boldsymbol{x}_0; \boldsymbol{\theta}_l)}{\sum_{l=1}^{L} p_l k(\boldsymbol{x}_0; \boldsymbol{\theta}_l)} \\
h(t|\boldsymbol{x}_0; G_L) &= \frac{f(t|\boldsymbol{x}_0; G_L)}{S(t|\boldsymbol{x}_0; G_L)} \\
m(t|\boldsymbol{x}_0; G_L) &= \frac{\int_t^\infty S(u|\boldsymbol{x}_0; G_L)du}{S(t|\boldsymbol{x}_0; G_L)} = \frac{E(t|\boldsymbol{x}_0; G_L) - \int_0^t S(u|\boldsymbol{x}_0; G_L)du}{S(t|\boldsymbol{x}_0; G_L)}
\end{aligned}
$$

where $K(t|\boldsymbol{x}_0;\boldsymbol{\theta}_l)$ is the conditional kernel distribution function (a gamma cdf under the product kernel of model (5)).

2.3 Simulation example

We simulate 1500 data values from a population having density: $f(t,x) = \sum_{l=1}^{M} q_l \Gamma(t; a_l, b_l) N(x; m_l, s_l^2)$, where $M = 6$, $a = (45, 3, 125, 0.4, 0.5, 4)'$, $b = (3, 0.2, 3.8, 0.2, 0.3, 5)'$, $m = (-12, -8, 0, 12, 18, 21)'$, $s = (6, 5, 4, 5, 3, 2)'$, and $q = (0.28, 0.1, 0.25, 0.21, 0.11, 0.05)'$. The simulated data is shown in the left panel of Figure 1.

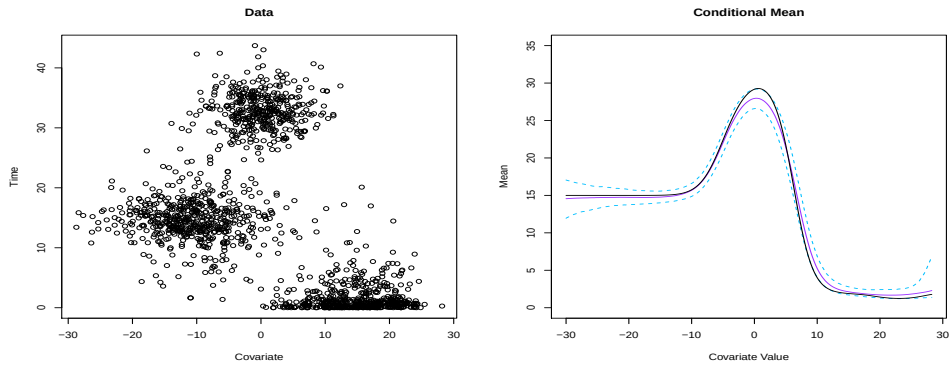

Figure 1: Simulated data (left), and point (purple solid) and interval estimates (light blue dashed) of the mean overlaying the truth (black solid) (right).

This population was constructed to have various shapes of the mrl function across different covariate values. The shapes include mrl functions with multiple change points, nonlinearly decreasing, UBT, and nonlinearly increasing. We demonstrate the ability of the joint DPMM in (4) to capture these various mrl functional forms at appropriate covariate values. The following priors were used: $a_\mu = (0.59, -2.12)$, $B_\mu = B_\sigma = ((0.019, 0.019)', (0.019, 0.019)')$, $a_\lambda = 0$, $a_\tau = 2$, $a_\rho = 1$, $b_\lambda = b_\tau = 88$, $b_\rho = 1/88$, $a_\alpha = 3$, $b_\alpha = 0.1$. The DP truncation level was set at $L = 80$.

The mean of the survival times across a grid of covariate values is shown in Figure 1 (right panel). In general, the model is able to capture the non-linear trend of the mean over the covariate values. The point estimate is almost on top of the truth for a large portion of the grid.

The results for mrl functional inference is shown in Figure 2. We provide point and interval estimates for the mrl function at six different covariate values. The model is able to capture the overall shape of the true mrl functions, despite the variety of shapes. At covariate values where the data is most dense, such as $x = -5$, $x = 0$, and $x = 5$, the inference is more precise as is seen in the narrow interval bands. As we move to covariate values away from zero, where data is more sparse, the wide interval bands reflect the uncertainty of the mrl functional shape.

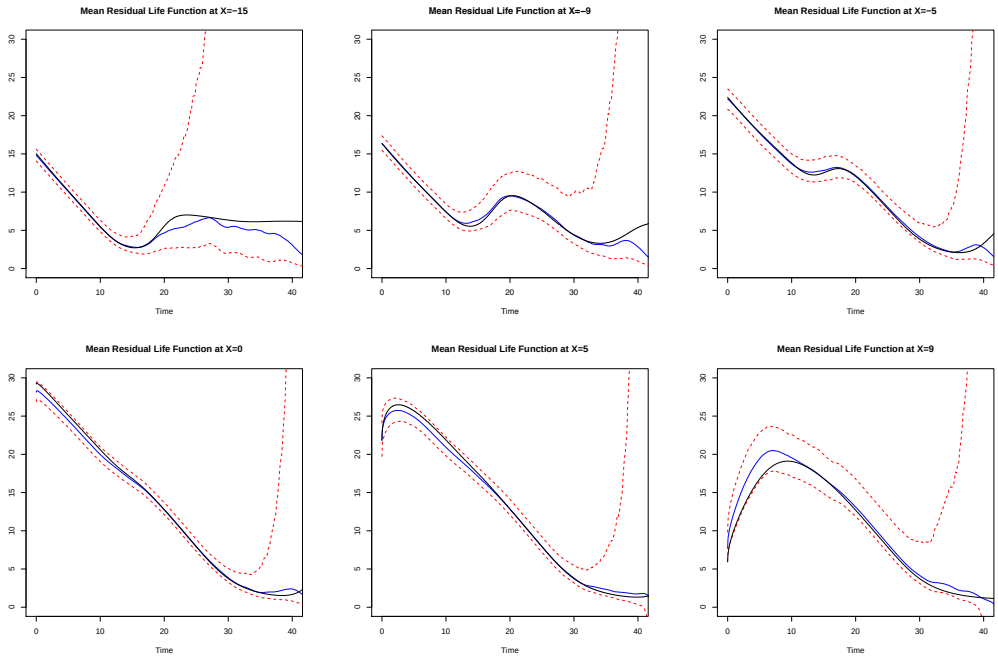
Mean Residual Life Function at X=−15
Mean Residual Life Function at X=−9
Mean Residual Life Function at X=−5
Mean Residual Life Function at X=0
Mean Residual Life Function at X=5
Mean Residual Life Function at X=9
Time

Figure 2: Point (blue solid) and 95% interval estimates (red dashed) of the mrl function for the specified covariate value overlaying the true mrl function of the population (black solid).

3 Dependent Dirichlet Process Mixture Model

Often in clinical trials, researchers are interested in modeling survival times of patients under treatment and control groups. Being that the underlining population pre treatment is typically the same, it is reasonable to expect that the survival distribution exhibit similarities across the experimental groups. The benefit of modeling the groups jointly is to be able to capture dependencies amongst groups, and to borrow strength from the group with a larger sample size in order to achieve more precise inference. Under this framework, we able to achieve nonstandard shapes in the mrl functions, that may even differ across groups contingent on the strength of the dependency across experimental groups.

3.1 Dirichlet process prior with dependent weights

Let $s \in S$ represent in general the index of dependence. In our case, this indicates the experimental group, that is $S = \{T, C\}$ where (T) is the treatment group and (C) is the control group. The regression DPMM in (2) can be extended to $f(t, \boldsymbol{x}; G_s) = \int_{\boldsymbol{\Theta}} k(t, \boldsymbol{x}; \boldsymbol{\theta}) dG_s(\boldsymbol{\theta})$ for $s \in S$, where now we are modeling a pair of dependent random mixing distributions $\{G_s : s \in S\}$. We desire to model the distributions in such a way as to incorporate dependencies across experimen-

tal groups, while maintaining marginally the DP prior, $G_s \sim DP$, for each $s \in S$. MacEachern (2000) develops the dependent DP prior in generality with both the weights and locations in the DP SB definition dependent on experimental group: $G_s(\cdot) = \sum_{l=1}^{\infty} \omega_{ls} \delta_{\boldsymbol{\theta}_{ls}}(\cdot)$. Marginally, $G_s \sim DP(\alpha_s, G_{0s})$ for each $s \in S$. MacEachern (2000) goes on to describe the computational difficulties in modeling dependencies in the weights across groups, thus motivating development of the "single p" model. In this model, the weights do not change over the groups, only the locations vary, $G_s(\cdot) = \sum_{l=1}^{\infty} \omega_l \delta_{\boldsymbol{\theta}_{ls}}(\cdot)$. We have studied such dependent DP mixture modeling for comparison of neuronal intensities under distinct experimental conditions in Kottas et al. (2012). Applications of "single p" dependent DP models include DeIorio et al. (2004), Gelfand et al. (2005), Rodriguez and ter Horst (2008), Kottas and Krnjajić (2009), DeIorio et al. (2009), and Fronczyk and Kottas (2010).

While computationally convenient and a useful extension of the basic DP prior, assuming the same weights has potential disadvantages in our setting. A practical disadvantage of the "single p" dependent DP construction involves applications with a moderate to large number of covariates. For such cases, the "single p" prior requires building dependence across $s \in S$ for a large number of kernel parameters, whereas modeling dependency through the weights is not affected by the dimensionality of the mixture kernel. In situations where we might expect similar locations across groups, modeling dependency through the weights is more attractive. In our context, we may expect the two groups to be comprised of similar components which however exhibit different prevalence across survival time.

We thus use mixing distributions of the form $G_s = \sum_{l=1}^{\infty} \omega_{ls} \delta_{\boldsymbol{\theta}_l}(\cdot)$, which under the truncation approximation, become $G_s \approx \sum_{l=1}^{L} p_{ls} \delta_{\boldsymbol{\theta}_l}$, for $s \in \{T, C\}$ representing the treatment and control groups, respectively. Hence, the dependent mixture model is given by

$$f(t, \boldsymbol{x}; G_s) \;=\; \int_{\boldsymbol{\Theta}} k(t, \boldsymbol{x}; \boldsymbol{\theta}) dG_s(\boldsymbol{\theta}) \approx \sum_{l=1}^{L} p_{ls} k(t, \boldsymbol{x}; \boldsymbol{\theta}_l) \quad \text{for } s \in \{T, C\}. \tag{6}$$

Under the stick-breaking method in obtaining the weights, we sample independently the latent parameters, $\upsilon_r \sim Beta(1, \alpha)$, which is equivalent to using $\zeta_r = (1 - \upsilon_r) \sim Beta(\alpha, 1)$ for $r = 1, ..., L-1$. If we use a bivariate beta distribution for (ζ_{Tr}, ζ_{Cr}), we can incorporate the dependency between the two groups. Minimally, we need the marginals to be $\zeta_{sr}^* \sim Beta(\alpha, 1)$ for $s \in \{T, C\}$. By applying a bivariate beta distribution to the dependent DPMM with common locations and dependent weights, the following model for survival regression data that specify

two experimental groups is presented,

$$\begin{aligned} t_{is}, \boldsymbol{x}_{is} | G_s &\overset{ind}{\sim} f(t_{is}, \boldsymbol{x}_{is}; G_s) = \int_{\boldsymbol{\Theta}} k(t_{is}, \boldsymbol{x}_{is}; \boldsymbol{\theta}) dG_s(\boldsymbol{\theta}), \quad i = 1, .., n_s \\ (G_s, G_{s'}) | \boldsymbol{\phi}, \boldsymbol{\psi} &\sim DDP(\boldsymbol{\phi}, G_0(\cdot; \boldsymbol{\psi})) \quad \text{for } s, s' \in \{T, C\}, s \neq s' \\ \boldsymbol{\theta}_l | \boldsymbol{\psi} &\overset{iid}{\sim} G_0(\cdot; \boldsymbol{\psi}), \quad l = 1, 2, ... \end{aligned}$$

where $G_s = \sum_{l=1}^{\infty} \omega_{ls} \delta_{\boldsymbol{\theta}_l} \approx \sum_{l=1}^{L} p_{ls} \delta_{\boldsymbol{\theta}_l}$, and the weights defined by $\omega_{1s} = 1 - \zeta_{1s}$, $\omega_{ls} = (1 - \zeta_{ls}) \prod_{r=1}^{l-1} \zeta_{rs}$ for $l \in \{2, 3, ...\}$ with $(\zeta_{lC}, \zeta_{lT}) | \boldsymbol{\phi} \overset{ind}{\sim}$ biv-beta$(\cdot | \boldsymbol{\phi})$ for $l \in \{1, 2, ...\}$. Marginally, $\zeta_{lC} \sim beta(\alpha_C, 1)$ and $\zeta_{lT} \sim beta(\alpha_T, 1)$. There are a number of bivariate beta distributions to consider, however, some exhibit more favorable properties for our purposes than others. In particular, more flexible marginals would allow for different α values, i.e., $\zeta_{Tr} \sim Beta(\alpha_T, 1)$ and $\zeta_{Cr} \sim Beta(\alpha_C, 1)$. Naturally, we also want the bivariate beta to correspond to reasonable computation of the MCMC, have a relatively simple density form, full support for the correlation between ζ_{Cr} and ζ_{Tr}, and ideally an analytic expression for the correlation. The correlation of the bivariate beta distribution for (ζ_{Tr}, ζ_{Cr}) will be important for the study of the implied dependence structure in the dependent DP prior for (G_T, G_C).

The bivariate beta distribution presented by Michael and Schucany (2011) has a simple analytic form for the correlation structure and while the density does not have an easily obtainable form, sampling from the density is straight forward. A bigger problem is that we can not obtain marginal distributions for ζ_{Cr} and ζ_{Tr} with different α parameters. Another possible bivariate beta is provided by Olkin and Liu (2003). This bivariate beta has a reasonable density form and the appropriate beta marginals for ζ_{Cr} and ζ_{Tr}, allowing different α values. However, the correlation does not have an analytic form, and the support of the correlation is restricted to positive values.

The bivariate beta that we chose to implement in our model is that of Nadarajah and Kotz (2005). They construct a bivariate beta distribution through products of beta distributions. Start with beta random variables, $U \sim Beta(a_1, b_1)$, $V \sim Beta(a_2, b_2)$, and $W \sim Beta(b, c)$, subject to the constraint, $c = a_1 + b_1 = a_2 + b_2$. The bivariate beta distribution we are interested in is defined for random variables, X and Y, where $X = UW$ and $Y = VW$. The marginals are given by $X \sim Beta(a_1, b_1 + b)$ and $Y \sim Beta(a_2, b_2 + b)$. We can obtain the desired beta marginals for ζ_{Cr} and ζ_{Tr} by setting $b_1 + b = b_2 + b = 1$, although the marginals will have the same α parameter. The density has a complicated form, but it can be sampled from using latent variables. The correlation has an analytic expression, however, has positive support, which may or may not be a reasonable assumption. Induced correlations in the model under this bivariate beta distribution is discussed in the next section.

The full hierarchical model, under this bivariate beta structure, can be written upon introducing the latent configuration variables, $\mathbf{w} = \{\mathrm{w}_{is} : i = 1, ..., n_s | s = C, T\}$, such that $w_{is} = l$ if the i^{th} observation at group s is assigned to mixture component l. Keeping the same kernel structure, baseline distribution, and priors as in (5), the hierarchical version of the model with the chosen bivariate beta distribution may be written as follows,

$$
\begin{aligned}
\{t_{is}\}|\mathbf{w}, \{\boldsymbol{\theta}_l\} &\sim \prod_{s\in\{C,T\}1} \prod_{i=1}^{n_s} \Gamma(t_{is}|e^{\theta_{\mathrm{w}_{is}}}, e^{\phi_{\mathrm{w}_{is}}}) \qquad (7)\\
\{x_{is}\}|\mathbf{w}, \{\boldsymbol{\theta}_l\} &\sim \prod_{s\in\{C,T\}1} \prod_{i=1}^{n_s} N(x_{is}|\beta_{\mathrm{w}_{is}}, \kappa^2_{\mathrm{w}_{is}})\\
\mathrm{w}_{is}|\{(\zeta_{ls})\} &\overset{ind}{\sim} \sum_{l=1}^{L}\{(1-\zeta_{ls})\prod_{r=1}^{l-1}\zeta_{rs}\}\delta_l(\mathrm{w}_{is}), \quad i = 1, ..., n_s \text{ and } s \in \{C, T\}\\
\{(\zeta_{lC}, \zeta_{lT})\}|\alpha, b &\sim Biv - Beta(\{(\zeta_{lC}, \zeta_{lT})\}|\alpha, b)\\
(\theta_l, \phi_l, \beta_l, \kappa^2_l)'|\boldsymbol{\mu}, \boldsymbol{\Sigma}, \lambda, \tau^2, \rho &\overset{iid}{\sim} N_2((\theta_l, \phi_l)'|\boldsymbol{\mu}, \boldsymbol{\Sigma})N(\beta_l|\lambda, \tau^2)\Gamma^{-1}(\kappa^2_l|a, \rho)
\end{aligned}
$$

where $\zeta_{lC} = UW$ and $\zeta_{lT} = VW$, for $l \in \{1, ..., L\}$, with $U \overset{iid}{\sim} Beta(\alpha, 1-b)$, $V \overset{iid}{\sim} Beta(\alpha, 1-b)$, and $W \overset{iid}{\sim} Beta(1+\alpha-b, b)$. We place the following priors: $\alpha \sim \Gamma(\alpha|a_\alpha, b_\alpha(\text{rate}))$, $b \sim Unif(0, 1)$, $\boldsymbol{\mu} \sim N_2(\boldsymbol{\mu}|a_\mu, B_\mu)$, $\boldsymbol{\Sigma} \sim IWish(\boldsymbol{\Sigma}|a_\Sigma, B_\Sigma)$, $\lambda \sim N(\lambda|a_\lambda, b_\lambda)$, $\tau^2 \sim \Gamma^{-1}(\tau^2|a_\tau, b_\tau)$, and $\rho \sim \Gamma(\rho|a_\rho, b_\rho)$. The posterior sampling algorithm details are provided in Appendix A.

3.2 Properties of the DDP mixture model

In what follows, we explore what correlation structures are induced by the Nadarajah and Kotz (2005) bivariate beta distribution. Under this bivariate beta construction, the groups have a common $\alpha = \alpha_C = \alpha_T$, and the correlation is driven by both parameters, α and b. The construction is based off of the product of independent beta distributions. Start with sampling the independent latent variables: $U \sim \beta(\alpha, 1-b)$, $V \sim \beta(\alpha, 1-b)$, $W \sim \beta(\alpha+1-b, b)$. Let $\zeta_C = UW$ and $\zeta_T = VW$. The weights are defined by $w_{s1} = 1-\zeta_{1s}$, $w_{ls} = (1-\zeta_{ls})\prod_{r=1}^{l-1}\zeta_{rs}$, for $l \in \{2, 3, ...\}$.

We are interested in obtaining the correlation between the two mixing distributions, G_C and G_T, implied under this bivariate beta distribution. We first start with the correlation between ζ_C and ζ_T, $Cor(\zeta_C, \zeta_T)$. We omit the component subscript in the latent variables, since results are the same for each $l \in \{1, 2, ...\}$. The covariance can be written as, $Cov(\zeta_C, \zeta_T) = E(\zeta_C\zeta_T) - E(\zeta_C)E(\zeta_T) = E((UW)(VW)) - E(UW)E(VW)$. Using the fact that U, V, W are independent, and that ζ_C and ζ_T have the same marginal distribution, $Beta(\alpha, 1)$, the covariance

and correlation reduces to:

$$\begin{aligned} Cov(\zeta_C, \zeta_T) &= \frac{\alpha^2 b}{(\alpha+1-b)(\alpha+1)^2(\alpha+2)} \\ Cor(\zeta_C, \zeta_T) &= \frac{\alpha b}{\alpha+1-b} \end{aligned} \tag{8}$$

The correlation between ζ_C and ζ_T can take values on the interval $(0, 1)$. As $b \to 0$ and/or $\alpha \to 0$, the correlation goes to 0. As $b \to 1$ and/or $\alpha \to \infty$, the correlation tends to 1.

The next step is to explore the correlation of the weights, $Cor(w_{lC}, w_{lT})$ for $l \in \{1, 2, ...\}$. When $l = 1$, $w_{1s} = 1 - \zeta_{1s}$, which is simply a linear operation, hence the covariance and correlation are the same as before. The $Cov(w_{1C}, w_{1T}) = Cov(\zeta_C, \zeta_T)$ and $Cor(w_{1C}, w_{1T}) = Cor(\zeta_C, \zeta_T)$ are given by (8). The case is different for $l = \{2, 3, ...\}$. In this case, the covariance is defined as $E\left[\left((1-\zeta_{lC})\prod_{r=1}^{l-1}\zeta_{rC}\right)\left((1-\zeta_{lT})\prod_{r=1}^{l-1}\zeta_{rT}\right)\right] - E\left[(1-\zeta_{lC})\prod_{r=1}^{l-1}\zeta_{rC}\right] E\left[(1-\zeta_{lT})\prod_{r=1}^{l-1}\zeta_{rT}\right]$. Using the fact that ζ_{ls} are independent across $l = 1, ..., L$, for each $s \in \{C, T\}$, the covariance, for $l \in \{2, 3, ...\}$, can be expressed as,

$$\begin{aligned} Cov(w_{lC}, w_{lT}) &= \frac{(\alpha+1-b)(\alpha+2)+\alpha^2 b}{(\alpha+1-b)(\alpha+1)^2(\alpha+2)} \left(\frac{\alpha^2 b + \alpha^2(\alpha+1-b)(\alpha+2)}{(\alpha+1-b)(\alpha+1)^2(\alpha+2)}\right)^{l-1} \\ &\quad - \frac{1}{(\alpha+1)^2}\left(\frac{\alpha^2}{(\alpha+1)^2}\right)^{l-1} \end{aligned} \tag{9}$$

The variance for the weights are independent of group, and can be expressed as $Var(w_{ls}) = 2/(\alpha+1)(\alpha+2)\left[\left(\alpha+\alpha^2(\alpha+2)\right)/\left((\alpha+1)^2(\alpha+2)\right)\right]^{l-1} - 1/(\alpha+1)^2\left[\alpha^2/(\alpha+1)^2\right]^{l-1}$. Therefore, the correlation, for $l \in \{2, 3, ...\}$, can be obtained by, $Cor(w_{lC}, w_{lT}) = Cov(w_{lC}, w_{lT})/Var(w_{ls})$, which is in closed form, but does not reduce. The correlation between the weights for $l \in \{2, 3, ...\}$ also takes values on the interval $(0, 1)$ and behaves the same in terms of the limits of α and b as in the case when $l = 1$. The component value, l, plays a slight role in the correlation, specifically as l get larger, the rate of change for smaller α values becomes less extreme.

We can now address the covariance and correlation between the two mixing distributions, $Cov(G_C, G_T)$ and $Cor(G_C, G_T)$. Let $B \in \boldsymbol{\Theta}$ represent a subset of the space of the mixing parameters. In the model we present, $\boldsymbol{\Theta}$ is equivalent to $\mathbb{R}^2$, so B is simply a subset of $\mathbb{R}^2$. Recall that the mixing distribution for group s has form $G_s(B) = \sum_{l=1}^{\infty} w_{ls}\delta_{\boldsymbol{\theta}_l}(B)$. Marginally, $G_s(B)$ follows a DP, so the expectation and variance of $G_s(B)$ is $G_0(B)$ and $G_0(B)[1 - G_0(B)]/(\alpha+1)$, respectively. The covariance between $G_C(B)$ and $G_T(B)$ is given by $Cov\left(\sum_{l=1}^{\infty} w_{lC}\delta_{\boldsymbol{\theta}_l}(B), \sum_{l=1}^{\infty} w_{lT}\delta_{\boldsymbol{\theta}_l}(B)\right)$, which boils down to the expression, $G_0(B)\sum_{l=1}^{\infty} w_{lC}w_{lT} + 2G_0^2(B)\sum_{l=1}^{\infty}\sum_{m=l+1}^{\infty} w_{lC}w_{mT} - G_0^2(B)$. The infinite series converges under geometric series (see appendix for details), and the covariance simplifies to be:

$$Cov(G_C(B), G_T(B)) = G_0(B)(1 - G_0(B))\left(\frac{(\alpha-2)b+\alpha+2}{\alpha(2\alpha-3b+5)-2b+2}\right) \tag{10}$$

The correlation, therefore, does not depend on the choice of B or G_0, it is driven by α and b alone:

$$Cor(G_C(B), G_T(B)) \quad = \quad \frac{(\alpha+1)((\alpha-2)b+\alpha+2)}{\alpha(2\alpha-3b+5)-2b+2} \tag{11}$$

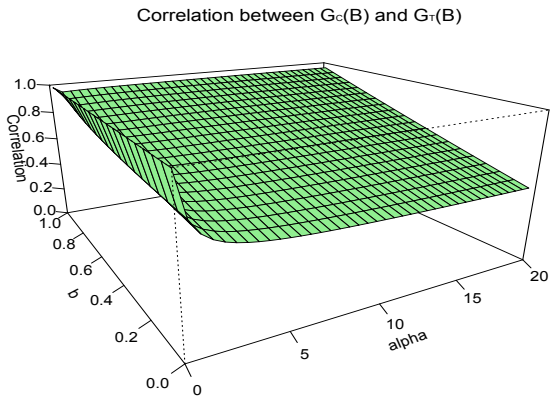

Figure 3: Correlation between $G_C(B)$ and $G_T(B)$ over a grid of α and b values.

The correlation of the mixing distribution lives on the interval $(1/2, 1)$, see Figure 3 for a visual. As $\alpha \to 0$ and/or $b \to 1$, the correlation tends to 1. When $\alpha \to \infty$ the correlation tends to $(b+1)/2$ and as $b \to 0$ the correlation tends to $(\alpha+1)/(2\alpha+1)$, so when $\alpha \to \infty$ and $b \to 0$ the correlation goes to $1/2$. Although this correlation space is limited, it is a typical range seen in the literature (e.g. McKenzie (1985)). It can easily be shown that the correlation of the survival distributions between the two groups given G_C and G_T also live on $(1/2, 1)$, which demonstrates the importance of prior knowledge of the relationship between the distributions of the two group survival times. While the possible values of correlation on the distributions of the survival times are restricted to $(1/2, 1)$, the correlation between the survival times across the two groups, $Cor(T_C, T_T)$, takes on values in $(0, 1)$. The correlation between T_C and T_T is found by marginalizing over the mixing distributions, G_C and G_T. Starting with the covariance, $Cov(T_C, T_T) = E[T_C T_T] - E[T_C]E[T_T] = E[E[T_C|G_C]E[T_T|G_T]] - E[E[T_C|G_C]]E[E[T_T|G_T]]$. Under the gamma kernel with bivariate normal G_0 that we have previously discussed, the covariance is given by the following,

$$Cov(T_C, T_T) \quad = \quad \left(e^{t_2'\boldsymbol{\mu}+\frac{1}{2}t_2'\boldsymbol{\Sigma}t_2} - e^{2(t_3'\boldsymbol{\mu}+\frac{1}{2}t_3'\boldsymbol{\Sigma}t_3}\right)\left(\frac{(\alpha-2)b+\alpha+2}{\alpha(2\alpha-3b+5)-2b+2}\right) \tag{12}$$

where $t_2 = (2, -2)'$ and $t_3 = (1, -1)'$. The variance of T_s, for both $s \in \{C, T\}$, is given by, $e^{t_1'\boldsymbol{\mu}+\frac{1}{2}t_1'\boldsymbol{\Sigma}t_1} + e^{t_2'\boldsymbol{\mu}+\frac{1}{2}t_2'\boldsymbol{\Sigma}t_2} - e^{2(t_3'\boldsymbol{\mu}+\frac{1}{2}t_3'\boldsymbol{\Sigma}t_3)}$. Recall that $t_1 = (1, -2)'$. Therefore the correlation is

given by,

$$Cor(T_C, T_T) = \left[\left(e^{t_2'\boldsymbol{\mu}+\frac{1}{2}t_2'\boldsymbol{\Sigma}t_2} - e^{2(t_3'\boldsymbol{\mu}+\frac{1}{2}t_3'\boldsymbol{\Sigma}t_3)}\right)\left(\frac{(\alpha-2)b+\alpha+2}{\alpha(2\alpha-3b+5)-2b+2}\right)\right] / \left[e^{t_1'\boldsymbol{\mu}+\frac{1}{2}t_1'\boldsymbol{\Sigma}t_1} + e^{t_2'\boldsymbol{\mu}+\frac{1}{2}t_2'\boldsymbol{\Sigma}t_2} - e^{2(t_3'\boldsymbol{\mu}+\frac{1}{2}t_3'\boldsymbol{\Sigma}t_3)}\right] \quad (13)$$

As the $e^{t_1'\boldsymbol{\mu}+\frac{1}{2}t_1'\boldsymbol{\Sigma}t_1} = E[e^{\theta-2\phi}] \to 0$ the correlation simplifies to $((\alpha-2)b+\alpha+2)/(\alpha(2\alpha-3b+5)-2b+2)$. In this case, as $\alpha \to 0$ the correlation tends to 1 and as $\alpha \to \infty$ the correlation tends to 0. Also, as $b \to 0$ the correlation tends to $1/(2\alpha+1)$ and as $b \to 1$ the correlation tends to $1/(\alpha+1)$. These results are scaled down as $E[e^{\theta-2\phi}]$, the expectation of the kernel variance, gets larger. In Figure 4, the surface plot of the correlation between the survival times are shown for different values of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are shown over a grid of values for b and α.

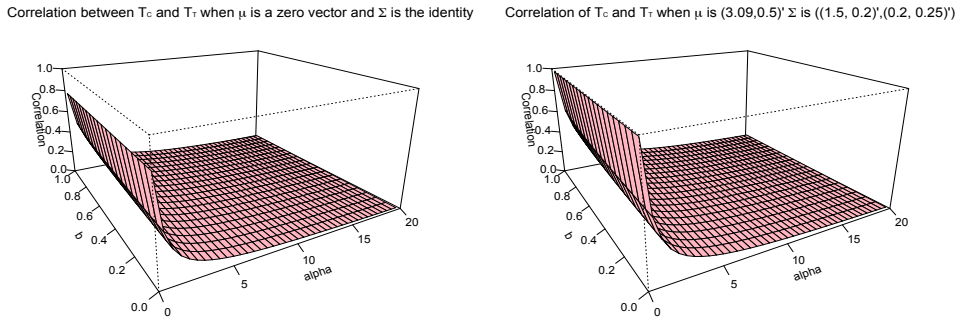
Correlation between Tc and TT when μ is a zero vector and Σ is the identity
Correlation of Tc and TT when μ is (3.09,0.5)' Σ is ((1.5, 0.2)',(0.2, 0.25)')
Correlation
b
alpha

Figure 4: Correlation between T_C and T_T when $\mu = (0,0)'$ and $\Sigma = ((1,0)'(0,1)')$ (left) and when $\mu = (3.09, 0.5)'$ and $\Sigma = ((1.5, 0.2)'(0.2, 0.25)')$ (right) over a grid of α and b values.

3.3 Simulation examples

In this section, we construct two sets of populations from which to sample from. The first set of populations is were constructed using a mixture of Weibull distributions that shared the same set of locations, but having different weights. Given that model (7) has DDP prior has the same construction, we would expect the model perform well under this situation. The populations for the first simulation is shown in the left panel in Figure 5. The panel shows how the two populations look similar having modes at the same locations, just differing prevalences for each mode. The second set of populations is also constructed using a mixture of Weibull distributions, however, this time we use different weights as well as locations. The intention is to test the model's inferential ability for populations that have quite different features. Figure 5 shows the density populations of the second simulation in the right panel. The second population

exhibits a single mode in between the two modes of the first population. The panel indicates the that the two densities are quite dissimilar.

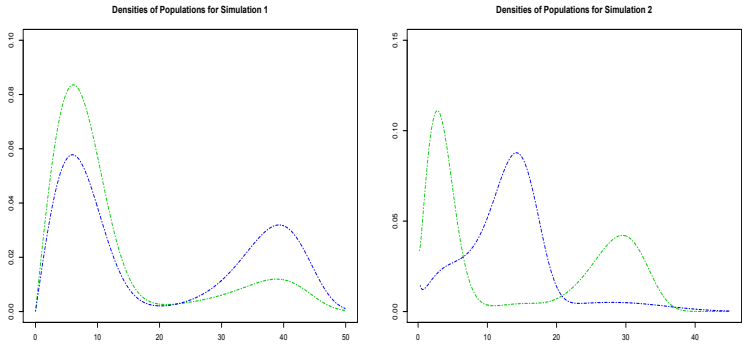

Figure 5: Simulation 1 population densities (left) and Simulation 2 population densities (right). The green curve represents the first population (T_1) while the blue represents the second (T_2) in each simulation.

3.3.1 Simulation 1

In Simulation 1, we demonstrate the model's ability to perform under circumstances in which resembles the structure of our model. Specifically, we simulate from two Weibull mixture distributions that share mixture locations, but have different weights:

$$T_1 \sim 0.7Weib(2,8) + 0.1Weib(3,10) + 0.05Weib(4,30) + 0.15Weib(8,40)$$
$$T_2 \sim 0.5Weib(2,8) + 0.05Weib(3,10) + 0.025Weib(4,30) + 0.425Weib(8,40)$$

The populations are comprised of four components each. We sample 250 survival times from the first population and 100 survival times from the second population. We do not consider censoring or covariates here. The histogram of the simulated survival data is shown in Figure 6

We place a Uniform prior on the b parameter and a Gamma prior on α with shape parameter 2 and rate parameter 0.8. The number of components is conservatively set at 40. Using prior specification methods discussed in Poynor and Kottas (2014), we place a bivariate normal prior on μ with mean vector $(1.87, 0.25)'$ and covariance matrix $((0.27, 0)', (0, 0.27)')$, and an inverse Wishart with 4 degrees of freedom and scale matrix $((0.27, 0)', (0, 0.27)')$. We update b and α together using a bivariate normal on the logit and log scale, respectively. The proposal is centered around the previous iteration, and initial MCMC runs are done to obtain an appropriate covariance matrix. After burn in and thinning, we obtain 2000 independent posterior samples.

Inference for the density, survival, and mrl functions are provided in Figure 7. The top panels are results for Group 1, while the bottom panels are that of Group 2. The colored solid

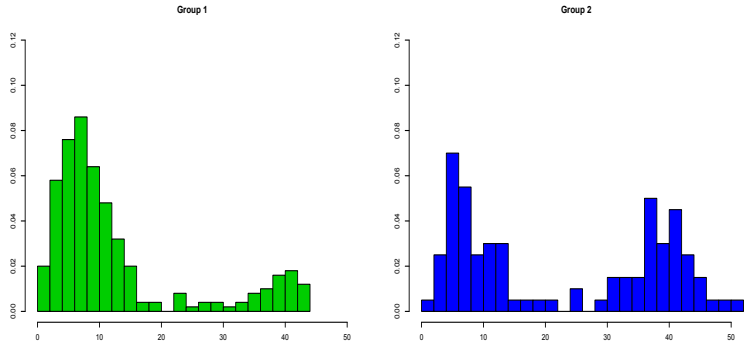

Figure 6: Simulation 1. Simulated survival times from mixture of Weibull having the same locations and different weights.

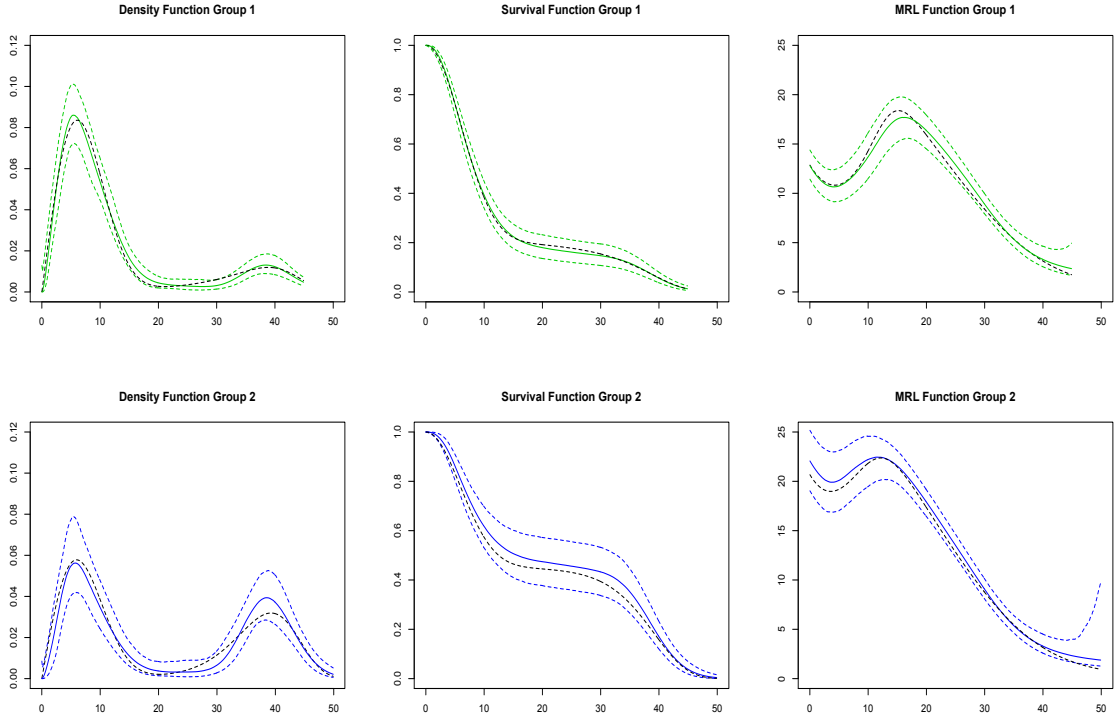
Density Function Group 1
Survival Function Group 1
MRL Function Group 1
Density Function Group 2
Survival Function Group 2
MRL Function Group 2

Figure 7: Simulation 1. Posterior point and 95% interval estimates for density (left), survival (middle), and mrl (right) functions. The truth is given by the black dashed curves.

and dashed curves represent the poster point and 95% interval estimates. The truth is plotted, in a dashed black curve, over the posterior results. The model is able to express the features of the functionals, and the true population density is captured within the 95% interval estimates. In particular, the flexibility of the model is demonstrated in the mrl function. The true mrl

is non-standard in both groups: initially decreasing, followed by an increase after about time 5, and then decreasing again after about time 12. The difference in sample size between the two groups is indicated by the slightly larger interval bands in Group 2 for the majority of the support of the data.

3.3.2 Simulation 2

The second simulation example is intended to be more of a challenge to the model. The populations consist of mixtures of Weibull distributions, however, here we use different weights, locations, and number of components. Group 1 is comprised of four components, while Group 2 is comprised of five:

$$T_1 \sim 0.5Weib(2,4) + 0.05Weib(0.6,4) + 0.025Weib(5,15) + 0.425Weib(8,30)$$
$$T_2 \sim 0.02Weib(0.6,1) + 0.02Weib(2,4) + 0.66Weib(5,15) + 0.2Weib(2,8) + 0.1Weib(4,30)$$

We simulate 250 observations from each population. All observations are fully observed, and no covariates are considered. The histogram of the simulated survival data is shown in Figure 8.

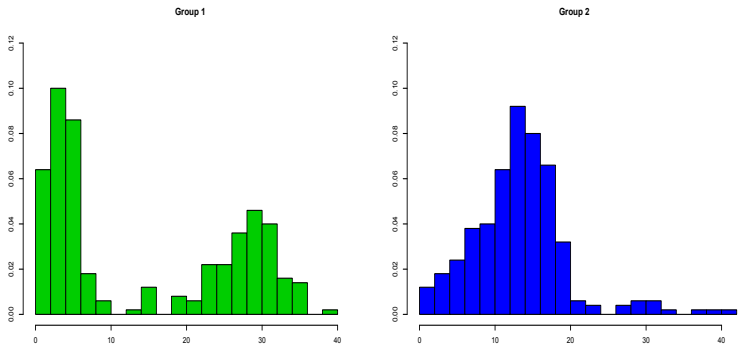

Figure 8: Simulation 2. Simulated survival times from mixture of Weibull having the same locations and different weights.

Once again, we use a uniform prior on b, and gamma prior on α with shape parameter 2 and rate parameter 0.8. The number of components is set at 40, which is a conservative value for these data. Using the same prior specification approach discussed in Poynor and Kottas (2014), a bivariate normal prior is placed on $\boldsymbol{\mu}$ with mean vector $(3.02, 0.54)'$ and covariance matrix $((0.1,0)', (0,0.1)')$. We update α and b the same way as in the first simulation. After burn in and thinning, we obtain 2000 independent posterior samples.

The posterior results for α, b, and the correlation between the mixing distributions are shown in Figure 10. The prior densities are shown in the plots as the red dashed line. The model favors smaller α values, which is not surprising since the number of components in the populations are small. The posterior for b also favors smaller values. Recall that in general, smaller b indicates smaller correlation. The model is likely trying to reflect the difference between the populations. On the other hand, smaller α values lead to a higher correlation. The posterior correlation between the mixing distributions seemw to settle between the competing values of α and b at around 0.7. The 95% credible interval for the correlation is given by $(0.619, 0.855)$.

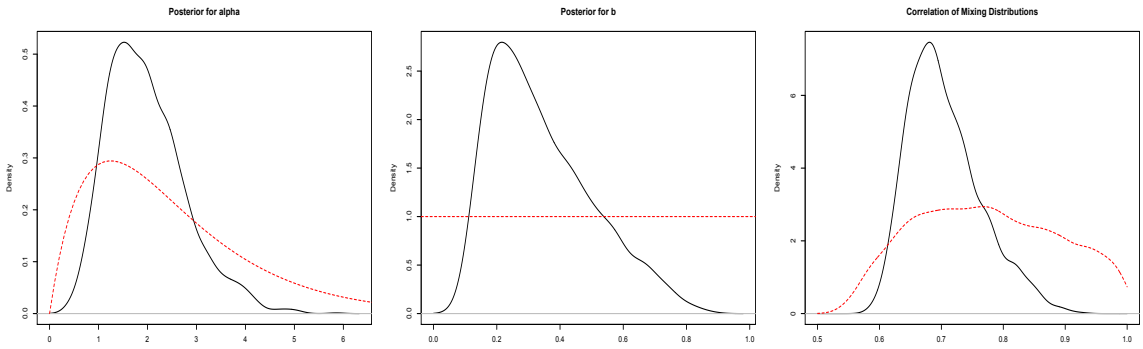

Figure 9: Simulation 2. Posterior point and 95% interval estimates for α (left), b (middle), and the correlation between the mixing distributions (right). The priors for α and b are given by the red dashed line.

The posterior results for the density, survival, and mrl functions are shown in Figure 10. Despite the difference in the features of the functionals between the two groups, the model is able to capture the features of each group with accuracy. This is especially exciting for the mrl functions. The mrl functions are quite different from one another, and both are non-standard shapes. The model has no problem capturing both shapes of the mrl functions. The only area where we can see struggle in the model for the mrl function inference is in the tails of the functionals. The true mrl function of Group 1 is slightly above the upper interval estimate of the model. This may be just due to the random nature of simulated data; this simulated data may suggest a lower mrl function in the tail. Another possibility is the extreme difference between the mrl functions of the two groups in the tails. Group 1 shoots up sharply, while Group 2 remains gradually decreasing. A third contributor to the tail struggle is that the sparsity of the data in this area, so models in general a have a tougher time achieving accuracy. Even with these elements against the model, the struggle is not significant.

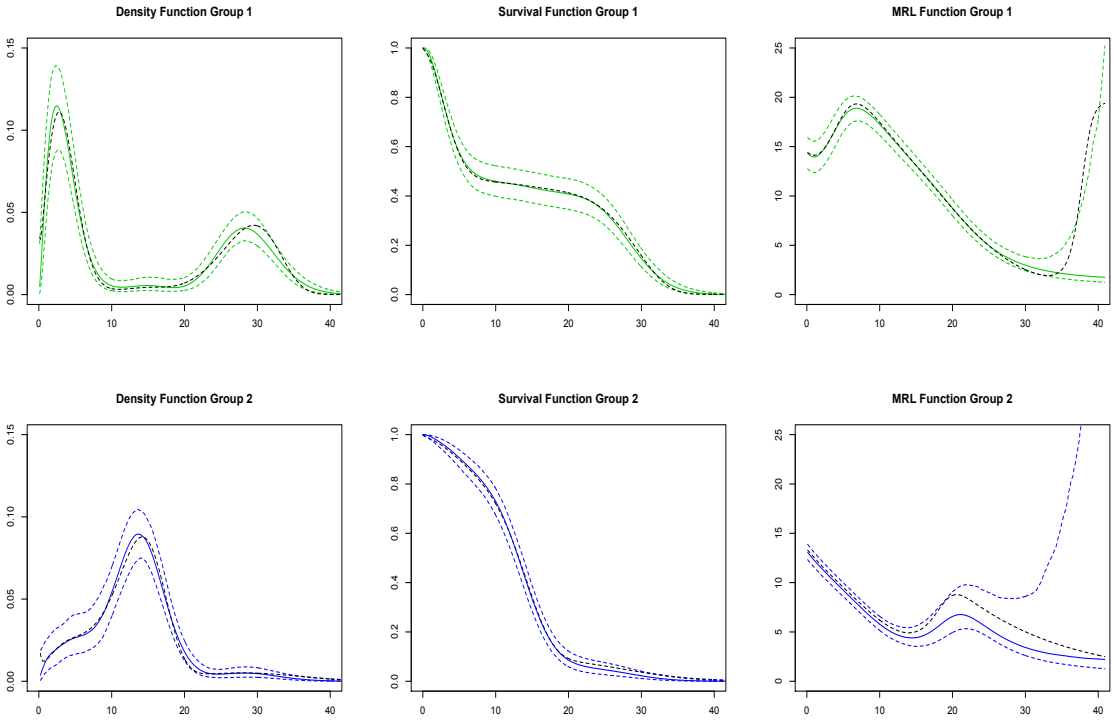
Density Function Group 1
Survival Function Group 1
MRL Function Group 1
Density Function Group 2
Survival Function Group 2
MRL Function Group 2

Figure 10: Simulation 2. Posterior point and 95% interval estimates for density (left), survival (middle), and mrl (right) functions. The truth is given by the black dashed curves.

The results from the two simulations demonstrates the utility of the gamma DDPMM. The model is able to incorporate dependency across two populations to achieve accurate inference in the functionals of each population. In particular, the model is able to capture provide flexible mrl inference for two groups that exhibit mrl functions having different features across the range of survival. In the following section, we apply the gamma DDPMM to a real dataset, and provide inferential results.

4 Small Cell Lung Cancer Data Example

In this section, we consider a dataset that describe the survival times, in days, of patients with small cell lung cancer under two treatment groups (Ying et al., 1995). The patients were randomly assigned to one of two treatments referred to as Arm A and Arm B. Arm A patients received cisplatin (P) followed by etoposide (E), while Arm B patients received (E) followed by (P). Arm A consists of 62 survival times, 15 of which are right censored. Arm B consists of 59 survival times, 8 of which are right censored. The age of each patient upon entry is also available,

however, in section 4.1, we will work with the treatment as the only covariate. We provide a brief discussion on the inferential results for these data under the gamma DDPMM developed in this paper with fitting the data under two independent gamma DP mixture models (DPMM) as was done in Poynor and Kottas (2014). We also compare formally the gamma DDPMM, to a handful of other models by computing the respective Conditional Predictive Ordinate (CPO) values. We also provide The age covariate is incorporated in section 4.2, and we compare the properties of the gamma DDPMM under the regression setting with the Linear DDP model developed in DeIorio et al. (2009).

4.1 Dependency across treatment groups

We fit a DDPMM using a gamma kernel to these data. Priors were specified using an analogous approach as described in Poynor and Kottas (2014), i.e., using the range and midrange of the observed survival times, which, in practice, would be specified by the expert. We place a uniform prior on b and a gamma prior with shape parameter 2 and rate parameter 0.5 is placed on α, and set $L = 80$. In Figure 11, the prior and posterior densities are overlaid for α, b, and the correlation between the mixing distributions, $Cor(G_C, G_T)$. The posterior densities for both α and b indicate learning for these parameters. Consequently, the model is able to learn about the correlation between the G_C and G_T. These data imply a fairly strong correlation between the mixing distributions as well as between the population distributions of the survival times under Arm A and Arm B.

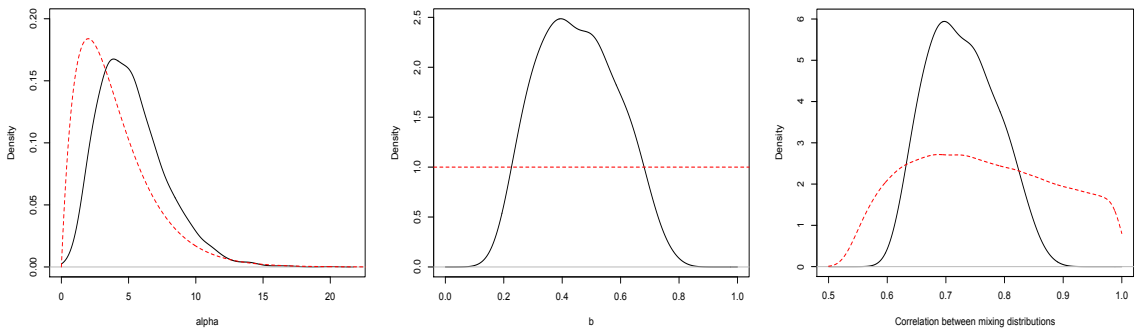

Figure 11: Prior (red dashed) and posterior (black solid) densities for α (left), b (middle), and $Cor(G_C, G_T)$ (right).

Inference for the density, survival, and mrl functions are provided in Figure 12. The point estimates for the density have the same general shape to the point estimates obtained by Kottas and Krnjajić (2009), who employ a semiparametric regression model. Both models indicate a

mode at about 450 days for Arm A and 350 days for Arm B. However, the point estimates under the gamma DDPMM are smoother than under the semi-parametric regression model for both groups. The difference is seen more obviously in the Arm B treatment. The point estimates for the two survival curves indicates that Arm A has a higher survival rate across the range of the data starting from about 200 days. When comparing the results under the gamma DDPMM from under the independent gamma DPMM, Poynor and Kottas (2014), the general conclusion regarding favorability of Arm A over Arm B remains the same, however, there an obvious change in the mrl functions. Although, the point estimates for the mrl functions maintain the same non-standard shape under both models, the separation between the two groups is far less under DDPMM compared to the DPMM (see, the figure 6 in Poynor and Kottas (2014)). Arm B is the group that appears to be most affected by the model change. Specifically, the point estimate for Arm B is shifted up. The shift is most drastic in the tail where data become more sparse.

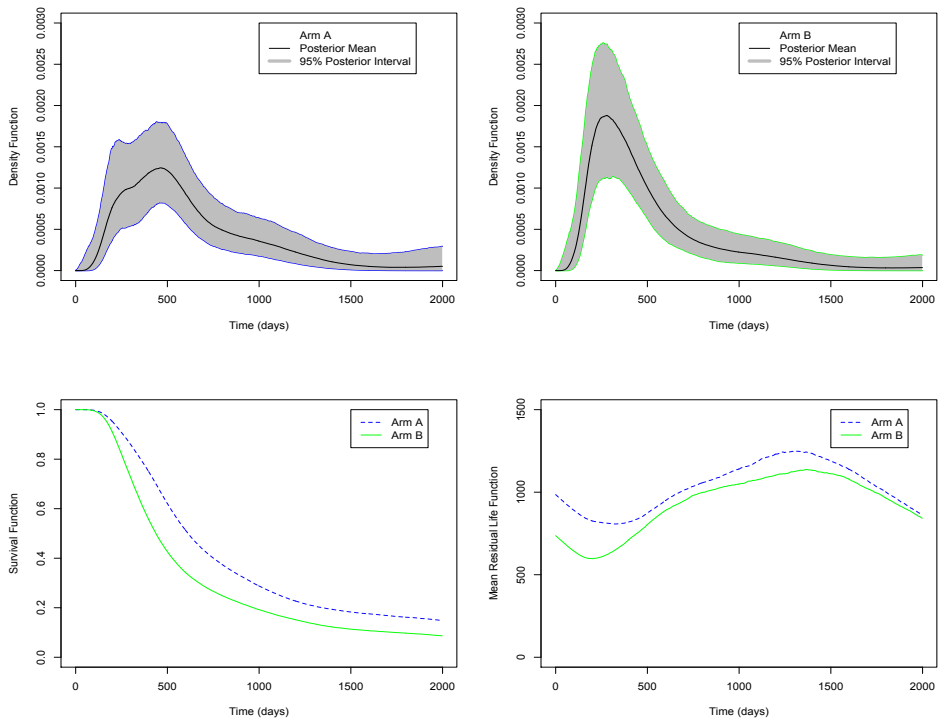

Figure 12: Posterior point and 95% interval estimates of the density function for Arm A (upper left) and Arm B (upper right). Posterior point estimate of the survival function (bottom left) and the mean residual life function (bottom right) for Arm A (blue dashed) and Arm B (green solid).

In Figure 13, we look at the prior probability, $Pr(m_A(t) > m_B(t))$, and posterior probability,

$Pr(m_A(t) > m_B(t)|data)$, under the gamma DDPMM. This Figure is analogous to figure 8 in Poynor and Kottas (2014). The prior probabilities under both models do not favor one mrl function over the other at any time point. We also see from the figure in Poynor and Kottas (2014) and Figure 13 here, that the posterior probability changes in a similar fashion as we move across the time space. Specifically, the probability is highest at smaller survival times then dips down followed by an increase and then then tapers back down. The range in probabilities is larger in Figure 13, with some probabilities reaching below 0.6. In particular, Figure 13 indicates a lower probability of the mrl function of Arm A being higher than the mrl function of Arm B after about 500 days when compared to the figure in Poynor and Kottas (2014).

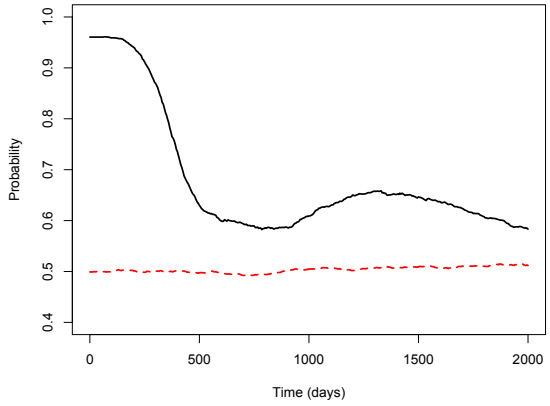

Figure 13: The posterior (black solid) and prior (red dashed) probability of the mrl function of Arm A being higher than the mrl function of Arm B over a grid of survival times (days).

We formally compare the performance of the gamma DDPMM to a parametric regression model that is able to obtain constant, increasing, decreasing, UBT, and BT shaped mrl functions. Namely, we extend the Exponentiated Weibull model (EWM) (Mudholkar and Strivasta, 1993) to a regression model through the scale parameter, σ, by setting $log(\sigma) = \beta_0 + \beta_1 x$. The survival function thus becomes, $S(t) = 1 - [1 - exp\left(-t^{\alpha} exp\left(\beta_0 + \beta_1 x\right)\right)]^{\theta}$. Here x takes value 0 when the survival time corresponds to Arm A, and a 1 when the survival time corresponds to Arm B. A normal prior with mean -10 and standard deviation 10 was placed on β_0, and a normal prior with mean0 and standard deviation 10 was placed on β_1. Exponential priors with rate parameters 1/1.1 and 1/0.9 were placed on α and θ, respectively. Prior parameters were specified by taking advantage of the closed form survival function of the EWM. The 10^{th}, 50^{th}, and 90^{th} quantiles of the data, which in practice would be specified by an expert, were used as arguments of the inverse survival function to obtain prior point estimates of the parameters.

The Conditional Predictive Ordinate (CPO), originally proposed as the leave-one-out cross-validation predictive density by Geisser and Eddy (1979), is a useful tool is assessing the performance of a model for a particular dataset. The CPO value of the i^{th} observation, CPO_i, is the marginal posterior probability of observing the i^{th} observation, t_i, when the model is fit to the data with t_i omitted:

$$CPO_i = f(t_i|data_{(-i)}) = \int f(t_i|\mathbf{\Psi}, x_i)\pi(\mathbf{\Psi}|data_{(-i)})d\mathbf{\Psi}$$

where $data_{(-i)}$ represents the data with the i^{th} observation removed, $\mathbf{\Psi}$ are the parameters of the model, x_i is the set of covariates associated with t_i, $f(\cdot)$ is the likelihood, and $\pi(\cdot)$ is the joint posterior distribution of the model parameters. A higher CPO value indicates a better model fit. A benefit of omitting the data value is that the data is only used once in assessing the model. The downfall, having to fit the model for each desired CPO value. However, CPO_i can be expressed in terms of the joint posterior distribution of the model parameters given ALL the observations:

$$CPO_i = \left(\int \frac{1}{f(t_i|\mathbf{\Psi}, x_i)}\pi(\mathbf{\Psi}|data)d\mathbf{\Psi}\right)^{-1}$$

In either case, the expression for CPO_i often does not have a closed form, so the MCMC approximation is used (see, for example, Chen et al. (2000)). The approximation must be monitored to ensure convergence is obtained. In some cases, the inverted likelihood can cause the estimator to be unstable. For these CPO values, the original definition must be used for the MCMC approximation. If the number of values of instability is small, this is not a large inconvenience, but if there are numerous instability cases, this model assessment technique may be practically unfeasible.

Obtaining the CPO values for the EWM is straight forward. The likelihood is given by $f(t_i|\alpha, \theta, \beta_0, \beta_1, x_i) = \theta\alpha t_i^{\alpha-1} exp\{\beta_0+\beta_1 x_i\}\left[1 - exp\left(-t_i^{\alpha} exp\{\beta_0 + \beta_1 x_i\}\right)\right]^{\theta-1}\left[exp\left(-t_i^{\alpha} exp\{\beta_0+\right.\right.$ $\left.\left.\beta_1 x_i\}\right)\right]$. Recall that x_i is 0 if t_i is an observation from Arm A and 1 if t_i is an observation from Arm B. Let M represent the number of MCMC iterations, then the CPO values for the EWM are estimated using the posterior samples of the parameters via the Harmonic mean estimator:

$$CPO_i \approx \left(\sum_{j=1}^{M} \frac{1}{f(t_i|\alpha_j, \theta_j, \beta_{0j}, \beta_{1j}, x_i)}\right)^{-1} \tag{14}$$

If the i^{th} observation is right censored, the likelihood is replaced by the survival function. The MCMC for the EWM was ran for 500000 iterations with a burn in of 10000 for an effective posterior sample size of 2000. The CPO values are plotted in red in Figure 14.

The DDPMM requires a slightly different expression for the CPO values. Under the DDPMM, each observation is believed to come from a particular component of the mixture, so while the density function is a mixture, the likelihood contribution of the i^{th} observation is a single component. Recall that the full posterior distribution from which we obtain samples from is $p(\boldsymbol{\theta}, \mathbf{w}, \boldsymbol{p}, \boldsymbol{\Psi})$, where $\boldsymbol{\Psi} = (\boldsymbol{\mu}, \boldsymbol{\Sigma})$. In the small cell lung cancer dataset, we have two experimental groups indexedby $s \in \{C, T\}$, thus for these data we compute CPO values for the i^{th} observation in group s, CPO_{is}. Again, let M represent the number of MCMC iterations, then the expression we need to approximate the CPO values using the posterior parameter samples is given by:

$$CPO_{is} \approx \frac{A}{B_{is}} \left(\sum_{j=1}^{M} \frac{\sum_{l=1}^{L} \Gamma(t_{is}|\boldsymbol{\theta}_{lj})}{\Gamma(t_{is}|\boldsymbol{\theta}_{\mathbf{w}_{isj}})} \right), \quad \text{where} \quad \frac{A}{B_{is}} = \left(\sum_{j=1}^{M} \frac{1}{\Gamma(t_{is}|\boldsymbol{\theta}_{\mathbf{w}_{isj}})} \right) \tag{15}$$

The details of this derivation is provided in Appendix B. The CPO values of the gamma DDPMM are plotted in blue (Arm A) and green (Arm B) in Figure 14.

With respect to the CPO values, the gamma DDPMM performs better than the EWM around the mode and the far tail of each density, but collectively it's a close race between the models. The shape of densities of each group are fairly standard, so it is not surprising that a flexible parametric model performs well. The EWM has higher CPO values in the tail of the fully observed data values, but struggles once the it enters the territory of the right censored values. This feature along with the lower CPO values around the mode is indicative of parametric nature of the EWM. Although the EWM performs well, it unable to capture the curvature of the mode while at the same time accurately estimating the skewness of the tail. An obvious compromise is made between capturing features of one aspect and not the other versus losing a little on both but not doing bad on either. The gamma DDPMM does not perform as well as the EWM in the tail of the observed data values, which may be attributed to the right censored observations that share some of the same time space as some of the last few fully observed survival times. A summary of the CPO values were obtained by averaging over the log of the CPO values, ALPML, in each group s:

$$ALPML_s = \tfrac{1}{n_s} \textstyle\sum_{i=1}^{n_s} log(CPO_{is})$$

The same rubric applies, meaning a higher ALPML indicates better model performance. The EWM scored a -6.09, while the gamma DDPMM scored a -6.05. The gamma DDPMM performs slightly better than the EWM. We are not too surprised the EWM scored so close to the gamma DDPMM being as the parametric model was selected after seeing previous inferential results for these data, and was chosen based off the posterior shape exhibited in the density.

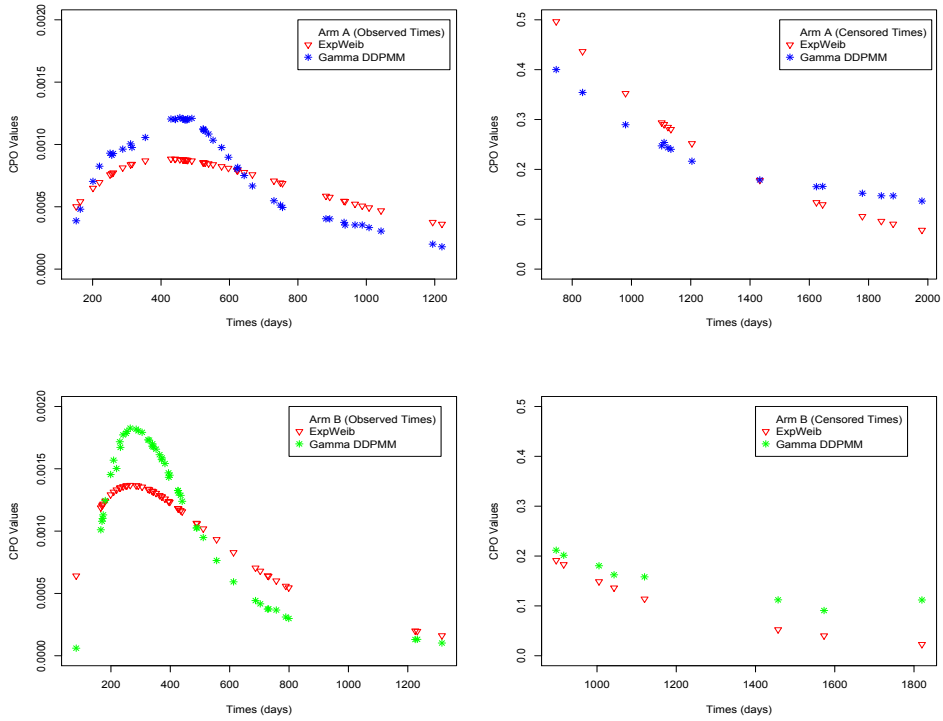

Figure 14: The CPO values under the EWM (red) and gamma DDPMM (blue and green) for the small cel lung cancer data. The top panels represent Arm A, and the bottom represent Arm B. The right column are the right censored survival times, and the left column are the fully observed survival times.

Kottas and Krnjajić (2009), provide a Bayesian semiparametric model for quantile regression that is a DP scale mixture of uniform densities. The model is fitted to the linear regression errors of the survival times, $\epsilon = t - \boldsymbol{x}^T\boldsymbol{\beta}$, assuming that the median and the mode of the error density is 0. They fit the model to the small cell lung cancer data, without the age covariate, and compute the CPO values. The ALPML that they reported is -6.91. Kottas and Krnjajić also consider a DP scale mixture of Laplace densities and a basic Weibull model for these data. Table 1 shows the ALPML value under each of the models.

Model	Summary Value
EWM	-6.09
DP scale mixture of uniform densities	-6.91
gamma DDPMM	-6.05
DP scale mixture of Laplace densities	-8.01
Weibull Model	-11.56

Table 1: Summary of the CPO values.

4.2 Incorporating the age covariate

Here, we incorporate the age (in years) of the subjects, upon entrance into the study, that is available to us in the small cell lung cancer dataset. The researchers did not select subjects from particular ages, so it is not a fixed covariate, and can be thought of as being random. Therefore, we model the age covariate on the log scale jointly with the survival response. Specifically, we used independent gamma and normal distributions for the survival times and age, respectively, see (4). We use prior specifications methods discussed in Poynor and Kottas (2014) for parameters associated with the survival times, and use a similar idea to specify parameters associated with the age covariate (see, e.g., Poynor (2010) for details). Appendix A show the details of the posterior sampling algorithm. We run the MCMC to obtain an effective posterior sample size of 2000.

In Figure 15, we plot the conditional mean of survival across a grid of ages. This inference was obtained by computing (3) at each posterior sample of the parameters for each experimental group. Specifically, the form of the mean regression for group $s \in \{A, B\}$ (A representing Arm A and B representing Arm B) at age x_0 under the gamma DDPMM, is given by,

$$E(t_s|x_0, G_s^L) = \sum_{l=1}^{L} \underbrace{p_{ls}\left(\frac{LN(x_0|\beta_l, \kappa_l^2)}{\sum_{r=1}^{L} p_{rs} LN(x_0|\beta_r, \kappa_r^2)}\right)}_{q_{ls}(x_0|\beta_l, \kappa_l^2)} e^{\theta_l - \phi_l} \tag{16}$$

Recall from Section 2, the set of functions $q_{ls}(\cdot)$ can be thought of as a new set of weights such that the mean regression is a finite weighted sum of the kernel component means. Moreover, the weights are functions of the covariate, indicating the potential of the model to capture non-standard relationships across the covariate space. This ability is demonstrated in Figure 15 where we an increase in the mean survival from about 36 to just after 50, followed by a steeper decline, particularly in Arm B, and then leveling out at higher ages.

We also look at the mrl function at age 50, 60, and 78, see Figure 16. At age 50, the mrl function for Arm A appear monotonic while the mrl of Arm B has a very shallow dip at about 400 days then becomes indistinguishable from Arm A. At age 60, the separation becomes more apparent towards in the earlier survival range, and the dips are more pronounced and present in both groups. At age 78, we see a similar curvature as in our past analysis: a dip around $300 - 400$ and a shallow mode around $1000 - 1200$. While the shapes and range of the mrl functions change across the covariate space, Arm A remains as high or higher than Arm B.

The linear DDP by DeIorio et al. (2009) is a leading model for Bayesian nonparametric survival regression. The linear DDP model works with a dependent DP structure for which the

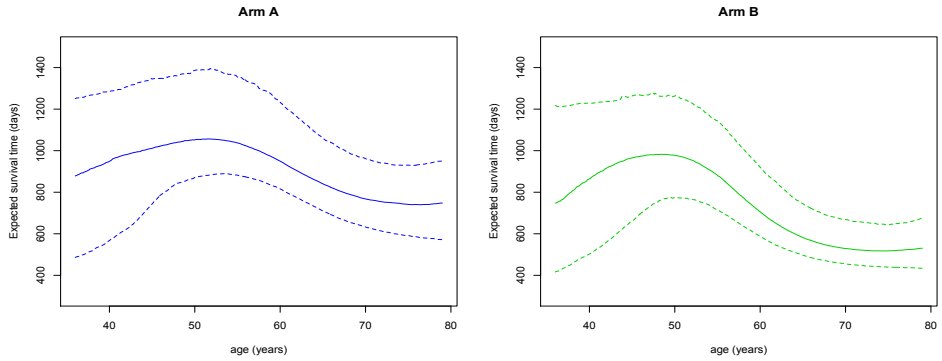

Figure 15: Point and 80% interval estimates of the conditional mean of the survival distribution of Arm A (blue) and Arm B (green) across a grid of age values (in years).

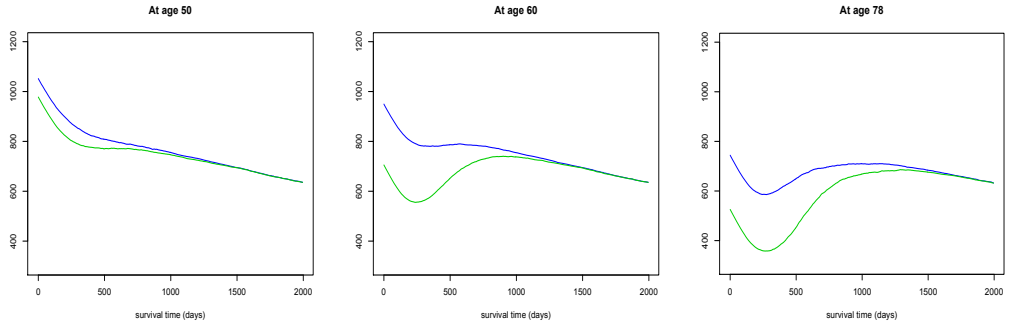

Figure 16: Estimates of the mrl function of Arm A (blue) and Arm B (green) for ages 50 (left), 60 (middle), and 78 (right), age is in years.

collection of random distributions, indexed by a set of covariates, x, denoted $\{G_x,\ x \in X\}$, is almost surely have the form of $G_x(\cdot) = \sum_{l=1}^{\infty} \omega_l \delta_{\theta_{xl}}(\cdot)$. Thus, the weights are shared across the covariate space while the locations differ. One of difference to note is that under the linear DDP, the covariates are not random, such that in the small cell lung cancer data set, the experimental group and the age would both be part of the set of covariates x. Specifically, under the linear DDP, if we let d_i be the design vector of the i^{th} survival time and α is the vector of coefficients, then $\theta_{x_i l} = \alpha_l d_i = m_l + A_{\nu l} + \beta_l z$. Here, m_l can be thought of the baseline effect, $A_{\nu l}$ represents the effect of a categorical covariate with outcome indicated by ν (experimental group indicated by s in our case), and β_l represents the effect of a continuous covariate z. They chose a normal kernel with mean αd_i and variance σ^2, mixing on both parameters. For $\mathbb{R}^+$ valued responses, modeling is performed on the log transformed data, which is the same as using a lognormal

kernel (Poynor, 2010). Thus the model is given by,

$$\begin{aligned} f_{x_i}(t_i|G) &\sim \int LN(t_i|\alpha d_i, \sigma^2) dG(\alpha, \sigma^2) \\ G &\sim DP(M, G_0) \end{aligned} \tag{17}$$

where we can write $G \sim DP(M, G_0)$ since $\theta = \alpha d_i$ implies $(\theta, \sigma^2) \sim G_{0x_i}$ by definition of base measure. Now, if we look at the implications of the mean regression function using the truncated version of the stick-breaking definition, we obtain (3.19) shown below.

$$E_x(t|G_L) = \sum_{l=1}^{L} p_l e^{(\alpha d + \sigma^2/2)} \tag{18}$$

The set of covariates only enters into the function as an exponential power, so the mean regression is a strictly increasing function in x (or equivalently d). While the linear DDP provides flexible inference for the survival and density functions, it may not be appropriate in case for which the regression mean or mrl is of interest. Even under a gamma kernel, which is a more appropriate kernel choice for mrl inference, reparameterizing to set the mean to αd would result in a regression mean function that is a weighted sum of linear functions: $E_x(t|G_L) = \sum_{l=1}^{L} p_l \alpha d$. Here, the relationship across the covariate space is subject to a linear trend. Therefore, under such circumstances, the gamma DDPMM would be preferable. In the discussion below, we describe extensions to include potential categorical covariates, however, if the number of experimental group exceeds two, a multivariate beta would have to be explored to extend the DDP.

5 Discussion

We have provided an example of a product kernel for a DPMM that jointly models survival times with a continuous covariate, and discussed the model in generality in the case of multiple continuous and discrete covariates. We provided a simulation study for a single continuous covariate to demonstrate the utility of the model. In particular, the model is able to achieve nonstandard shapes in the mrl functions across the covariate space. We also developed a DDP prior using a bivariate beta distribution to model dependency between two experimental groups. Combing these two modeling techniques, we are able to simultaneously incorporate information from both experimental groups as well as associated covariates, reducing uncertainty while still maintaining flexibility in the inferential results.

APPENDIX A: Posterior sampling from the gamma DDPMM with random covariates

Here we show the algorithm used to obtain the posterior samples of the parameters in a gamma DDPMM is the presence of a single random continuous real-valued covariate. We use the blocked Gibbs sampler described in Ishwaran and James (2001), with Metropolis-Hastings steps when conjugacy is not obtainable. The full hierarchical version of the model is written as,

$$\begin{aligned}
(t_{is}, x_{is})|\mathrm{w}_{is}, \boldsymbol{\theta}_l &\overset{ind}{\sim} \Gamma(t_{is}|e^{\theta_{\mathrm{w}_{is}}}, e^{\phi_{\mathrm{w}_{is}}})N(x_{is}|\beta_{\mathrm{w}_{is}}, \kappa^2_{\mathrm{w}_{is}}) \\
\mathrm{w}_{is}|\{(\zeta_{ls})\} &\overset{ind}{\sim} \sum_{l=1}^{L}\{(1-\zeta_{ls})\prod_{r=1}^{l-1}\zeta_{rs}\}\delta_l(\mathrm{w}_{is}), \quad \text{for } i = 1, ..., n_s \text{ and } s \in \{C, T\} \\
\{(\zeta_{lC}, \zeta_{lT})\}|\alpha, b &\sim Biv - Beta(\{(\zeta_{lC}, \zeta_{lT})\}|\alpha, b) \\
(\theta_l, \phi_l)'|\boldsymbol{\mu}, \boldsymbol{\Sigma} &\overset{iid}{\sim} N_2((\theta_l, \phi_l)'|\boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad \text{for } l = 1, ..., L
\end{aligned}$$

where $\zeta_{lC} = UW$ and $\zeta_{lT} = VW$ for $l \in \{1, ..., L\}$ with $U \overset{iid}{\sim} Beta(\alpha, 1-b)$, $V \overset{iid}{\sim} Beta(\alpha, 1-b)$, and $W \overset{iid}{\sim} Beta(1+\alpha-b, b)$. We place the following priors: $\alpha \sim \Gamma(\alpha|a_\alpha, b_\alpha)$, $b \sim Unif(b|0,1)$, $\boldsymbol{\mu} \sim N_2(\boldsymbol{\mu}|a_\mu, B_\mu)$, $\boldsymbol{\Sigma} \sim IWish(\boldsymbol{\Sigma}|a_\Sigma, B_\Sigma)$, $\beta_l|\lambda, \tau^2 \overset{iid}{\sim} N(\beta_l|\lambda, \tau^2)$, $\kappa_l^2|a$, $\rho \overset{iid}{\sim} \Gamma^{-1}(\kappa^2|a, \rho)$, $\lambda \sim N(\lambda|a_\lambda, b_\lambda^2)$, $\tau^2 \sim \Gamma^{-1}(\tau^2|a_\tau, b_\tau)$, and $\rho \sim \Gamma(\rho, a_\rho, b_\rho)$, with $l \in \{1, ..., L\}$. Let L_s^* be the number of distinct components, and $\mathbf{w}_s^* \equiv \{\mathrm{w}_{js}^* : j = 1, ..., L_s^*\}$ be the vector of distinct components for group $s \in \{C, T\}$. For $i = 1, ..., n_s$, let $\delta_{is} = 0$ if t_{is} is observed and $\delta_{is} = 1$ if t_{is} is right censored for $s \in \{C, T\}$. Let Ψ represent the vector of the most recent iteration of all other parameters. Let $b = 1, ..., B$ be the number of iterations in the MCMC. The posterior samples of $p(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\beta}, \boldsymbol{\kappa}^2, \mathbf{w}, \boldsymbol{\zeta}, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \lambda, \tau^2, \rho, \alpha, b|data)$ can be obtained by the following algorithm:

Sample from the posterior conditional distribution for $(\theta_l, \phi_l)'$,β_l, and κ_l^2 for $l = 1, ..., L$: If l is not already a component: $l \notin \mathbf{w}_C^{*(b)} \cup \mathbf{w}_T^{*(b)}$

$$\begin{aligned}
p(\theta_l^{(b+1)}, \phi_l^{(b+1)}|data, \Psi) &\overset{\text{draw}}{\sim} N_2(\boldsymbol{\mu}^{(b)}, \boldsymbol{\Sigma}^{(b)}) \\
p(\beta_l^{(b+1)}|data, \Psi) &\overset{\text{draw}}{\sim} N(\lambda^{(b)}, \kappa_l^{2(b)}) \\
p(\kappa_l^{2(b+1)}|data, \Psi) &\overset{\text{draw}}{\sim} \Gamma^{-1}(a, \rho^{(b)})
\end{aligned}$$

If l is an active component in either or both: $l \in \mathbf{w}_C^{*(b)} \cup l \in \mathbf{w}_T^{*(b)}$. Then $p(\theta_l, \phi_l|data, \Psi) \propto N_2((\theta_l, \phi_l)'|\boldsymbol{\mu}, \boldsymbol{\Sigma}) \prod_{s\in\{C,T\}} \prod_{\{i:l=\mathrm{w}_{is}\}} \left[\Gamma(t_{is}|e^{\theta_l}, e^{\phi_l})\right]^{1-\delta_{is}} \left[\int_{t_{is}}^{\infty} \Gamma(u_i|e^{\theta_l}, e^{\phi_l})dt_i\right]^{\delta_{is}}$.
We use a Metropolis-Hastings step for this update. We sample from the proposal distribution $(\theta_l', \phi_l')' \sim N_2((\theta_l^{(b)}, \phi_l^{(b)})', cS^2)$, where S^2 is updated from the average posterior samples of $\boldsymbol{\Sigma}$ under initial runs, and $c > 1$. Draw $\eta \sim Unif(0, 1)$.

If $\eta < min\{1,$

$$\left.\frac{N_2((\theta_l',\phi_l')'|\boldsymbol{\mu}^{(b)},\boldsymbol{\Sigma}^{(b)})\prod_{s\in\{C,T\}}\prod_{\{i:l=\mathrm{w}_{is}^{(b)}\}}\left[\Gamma(t_{is}|e^{\theta_l'},e^{\phi_l'})\right]^{1-\delta_{is}}\left[\int_{t_{is}}^{\infty}\Gamma(u_{is}|e^{\theta_l'},e^{\phi_l'})dt_i\right]^{\delta_{is}}}{N_2((\theta_l^{(b)},\phi_l^{(b)})'|\boldsymbol{\mu}^{(b)},\boldsymbol{\Sigma}^{(b)})\prod_{s\in\{C,T\}}\prod_{\{i:l=w_{is}^{(b)}\}}\left[\Gamma(t_{is}|e^{\theta_l^{(b)}},e^{\phi_l^{(b)}})\right]^{1-\delta_{is}}\left[\int_{t_{is}}^{\infty}\Gamma(u_{is}|e^{\theta_l^{(b)}},e^{\phi_l^{(b)}})dt_i\right]^{\delta_{is}}}\right\}$$

set $(\theta_l^{(b+1)},\phi_l^{(b+1)})' = (\theta_l',\phi_l')'$
else $(\theta_l^{(b+1)},\phi_l^{(b+1)})' = (\theta_l^{(b)},\phi_l^{(b)})'$.
For β_l and κ_l, we have $p(\beta_l|data,\Psi) \propto N(\beta_l|\lambda,\tau^2)\prod_{s\in\{C,T\}}\prod_{\{i:l=\mathrm{w}_{is}\}}N(x_{is}|\beta_l,\kappa_l^2)$
and $p(\kappa_l^2|data,\Psi) \propto \Gamma^{-1}(\kappa_l^2|a,\rho)\prod_{s\in\{C,T\}}\prod_{\{i:l=\mathrm{w}_{is}\}}N(x_{is}|\beta_l,\kappa_l^2)$. Thus we sample via:

$$\begin{aligned}
p(\beta_l^{(b+1)}|data,\Psi) &\overset{\text{draw}}{\sim} N(m_\beta,s_\beta^2)\\
p(\kappa_l^{2(b+1)}|data,\Psi) &\overset{\text{draw}}{\sim} \Gamma^{-1}\left(0.5\left[\sum_{s\in\{C,T\}}\sum_{\{i:l=\mathrm{w}_{is}\}}1\right]+a, 0.5\left[\sum_{s\in\{C,T\}}\sum_{\{i:l=\mathrm{w}_{is}\}}(x_{is}-\beta_l^{(b+1)})^2\right]+\rho^{(b)}\right)
\end{aligned}$$

where $m_\beta = s_\beta^2\left(\kappa_l^{-2(b)}\left[\sum_{s\in\{C,T\}}\sum_{\{i:l=\mathrm{w}_{is}\}}x_{is}\right]+\tau^{-2(b)}\lambda^{(b)}\right)$,
and $s_\beta^2 = \left(\tau^{-2(b)}+\kappa_l^{-2(b)}\left[\sum_{s\in\{C,T\}}\sum_{\{i:l=\mathrm{w}_{is}\}}1\right]\right)^{-1}$.

To obtain samples from $p(\boldsymbol{\zeta}|\Psi.data)$ we work with the latent variables $\{U_l,V_l,W_l\}$. We have $p(\{(U_l,V_l,W_l)\}|\Psi,data) \propto \prod_{l=1}^{L-1}U_l^{(\sum_{r=l+1}^{L}M_{rC})+\alpha-1}(1-U_l)^{-b}(1-U_lW_l)^{M_{lC}}V_l^{(\sum_{r=l+1}^{L}M_{rT})+\alpha-1}(1-V_l)^{-b}(1-V_lW_l)^{M_{lT}}W_l^{(\sum_{r=l+1}^{L}M_{rC}+M_{rT})+\alpha-b}(1-W_l)^{b-1}$. Using slice sampling, we can introduce latent variables ν_l and γ_l for $l=1,...,L$, such that we have Gibbs steps for for each parameter. The joint posterior of interest becomes: $p(\{(U_l,V_l,W_l,\nu_l,\gamma_l)\}|\Psi,data) \propto \prod_{l=1}^{L-1}U_l^{(\sum_{r=l+1}^{L}M_{rC})+\alpha-1}(1-U_l)^{-b}\mathbf{1}_{\left(0<\nu_l\le(1-U_lW_l)^{M_{lC}}\right)}V_l^{(\sum_{r=l+1}^{L}M_{rT})+\alpha-1}(1-V_l)^{-b}\mathbf{1}_{\left(0<\gamma_l\le(1-V_lW_l)^{M_{lT}}\right)}W_l^{(\sum_{r=l+1}^{L}M_{rC}+M_{rT})+\alpha-b}(1-W_l)^{b-1}$ Therefore, we have the following Gibbs steps for $l=1,...,L-1$

$$\begin{aligned}
p(\nu_l^{(b+1)}|\Psi,data) &\sim Unif\left(0,(1-U_l^{(b)}W_l^{(b)})^{M_{lC}^{(b)}}\right)\\
p(\gamma_l^{(b+1)}|\Psi,data) &\sim Unif\left(0,(1-V_l^{(b)}W_l^{(b)})^{M_{lT}^{(b)}}\right)\\
p(U_l^{(b+1)}|\Psi,data) &\sim Beta\left(\left(\sum_{r=l+1}^{L}M_{rC}^{(b)}\right)+\alpha,1-b\right)\mathbf{1}_{\left(0,\frac{1}{W_l^{(b)}}\left[1-exp\left(\frac{log(\nu_l^{(b+1)})}{M_{lC}^{(b)}}\right)\right]\right)}\\
p(V_l^{(b+1)}|\Psi,data) &\sim Beta\left(\left(\sum_{r=l+1}^{L}M_{rT}^{(b)}\right)+\alpha,1-b\right)\mathbf{1}_{\left(0,\frac{1}{W_l^{(b)}}\left[1-exp\left(\frac{log(\gamma_l^{(b+1)})}{M_{lT}^{(b)}}\right)\right]\right)}\\
p(W_l^{(b+1)}|\Psi,data) &\sim Beta\left(\left(\sum_{r=l+1}^{L}M_{rT}^{(b)}+M_{rC}^{(b)}\right)+\alpha+1-b,b\right)\mathbf{1}_{(0,m^*)}
\end{aligned}$$

where $m^* = min\left\{\frac{1}{U_l^{(b+1)}}\left[1-exp\left(\frac{log(\nu_l^{(b+1)})}{M_{lT}^{(b)}}\right)\right], \frac{1}{V_l^{(b+1)}}\left[1-exp\left(\frac{log(\gamma_l^{(b+1)})}{M_{lT}^{(b)}}\right)\right]\right\}$
Set $\zeta_{lC}^{(b+1)} = U_l^{(b+1)}W_l^{(b+1)}$ and $\zeta_{lT}^{(b+1)} = V_l^{(b+1)}W_l^{(b+1)}$

For the update for w_{is} for $i = 1, ..., n_s$ and $s \in \{C, T\}$ we have $p(\mathrm{w}_{is}|data, \Psi) \propto \Gamma(t_{is}|e^{\theta_{\mathrm{w}_{is}}}, e^{\phi_{\mathrm{w}_{is}}})N(x_{is}|\beta_{\mathrm{w}_{is}}, \kappa^2_{\mathrm{w}_{is}}) \sum_{l=1}^{L}\{(1-\zeta_{ls})\prod_{r=1}^{l-1}\zeta_{rs}\}\delta_l(\mathrm{w}_{is})$, so we sample:

$$p(\mathrm{w}_{is}^{(b+1)}|data, \Psi) \overset{\text{draw}}{\sim} \sum_{l=1}^{L}\tilde{p}_{lis}\delta_{(l)}(\mathrm{w}_{is})$$

where $\tilde{p}_{lis} = \left\{p_{ls}\left[\Gamma(t_{is}|e^{\theta_l^{(b+1)}}, e^{\phi_l^{(b+1)}})\right]^{1-\delta_{is}}\left[\int_{t_{is}}^{\infty}\Gamma(u_{is}|e^{\theta_l^{(b+1)}}, e^{\phi_l^{(b+1)}})du_{is}\right]^{\delta_{is}}N(x_{is}|\beta_l^{(b+1)}, \kappa_l^{2(b+1)})\right\} / \left\{\sum_{l=1}^{L}p_{ls}\left[\Gamma(t_{is}|e^{\theta_l^{(b+1)}}, e^{\phi_l^{(b+1)}})\right]^{1-\delta_{is}}\left[\int_{t_{is}}^{\infty}\Gamma(u_{is}|e^{\theta_l^{(b+1)}}, e^{\phi_l^{(b+1)}})du_{is}\right]^{\delta_{is}}N(x_{is}|\beta_l^{(b+1)}, \kappa_l^{2(b+1)})\right\}$ with $p_{1s} = 1-\zeta_{1s}$ and $p_{ls} = (1-\zeta_{ls})\prod_{r=1}^{l-1}\zeta_{rs}$ for $l = 2, ..., L-1$.

For the update for $\boldsymbol{\mu}$ we have $p(\boldsymbol{\mu}|data, \Psi) \propto N_2(\boldsymbol{\mu}|a_\mu, B_\mu)\prod_{l=1}^{L}N_2((\theta_l, \phi_l)'|\boldsymbol{\mu}, \boldsymbol{\Sigma})$, so we sample:

$$p(\boldsymbol{\mu}^{(b)}|data, \Psi) \overset{\text{draw}}{\sim} N_2(m_\mu, S_\mu^2)$$

where $m_\mu = S_\mu^2\left(B_\mu^{-1}a_\mu + \boldsymbol{\Sigma}^{-1}\sum_{l=1}^{L}\boldsymbol{\theta}_l^{(b)}\right)$, $S_\mu^2 = \left(B_\mu^{-1} + L\boldsymbol{\Sigma}^{-1(b)}\right)^{-1}$. Turning to the update of $\boldsymbol{\Sigma}$, we have $p(\boldsymbol{\Sigma}|data, \Psi) \propto \prod_{l=1}^{L}N_2((\theta_l, \phi_l)'|\boldsymbol{\mu}, \boldsymbol{\Sigma})IWish(\boldsymbol{\Sigma}|a_\Sigma, B_\Sigma)$, so we sample:

$$p(\boldsymbol{\Sigma}^{(b+1)}|data, \Psi) \overset{\text{draw}}{\sim} IWish(L + a_\Sigma, B_\Sigma + \sum_{l=1}^{L}(\boldsymbol{\theta}_l^{(b+1)} - \boldsymbol{\mu}^{(b+1)})(\boldsymbol{\theta}_l^{(b+1)} - \boldsymbol{\mu}^{(b+1)})')$$

For the update for λ we have $p(\lambda|data, \Psi) \propto N(\lambda|a_\lambda, b_\lambda^2)\prod_{l=1}^{L}N(\beta_l|\lambda, \tau^2)$, so we sample:

$$p(\lambda^{(b+1)}|data, \Psi) \overset{\text{draw}}{\sim} N(m_\lambda, s_\lambda^2)$$

where $m_\lambda = s_\lambda^2\left(b_\lambda^{-2}a_\lambda + \tau^{-2}\sum_{l=1}^{L}\beta_l\right)$ and $s_\lambda^2 = \left(b_\lambda^{-2} + \tau^{-2(b)}L\right)^{-1}$. For the update for τ^2 we have $p(\tau^2|data, \Psi) \propto \Gamma^{-1}(\tau^2|a_\tau, b_\tau)\prod_{l=1}^{L}N(\beta_l|\lambda, \tau^2)$, so we sample:

$$p(\tau^{2(b+1)}|data, \Psi) \overset{\text{draw}}{\sim} \Gamma^{-1}\left(0.5L + a_\tau, 0.5\left[\sum_{l=1}^{L}(\beta_l^{(b+1)} - \lambda^{(b+1)})^2\right] + b^\tau\right)$$

For the update for ρ, $p(\rho|data, \Psi) \propto \Gamma(\rho|a_\rho, b_\rho)\prod_{l=1}^{L}\Gamma^{-1}(\kappa_l^2|a, \rho)$, so we sample:

$$p(\rho^{(b+1)}|data, \Psi) \overset{\text{draw}}{\sim} \Gamma\left(aL + a_\rho, \left[\sum_{l=1}^{L}\kappa_l^{-2(b+1)}\right] + b_\rho\right)$$

We do not have conjugacy for α and b, so we turn to the Metropolis-Hastings algorithm to update these parameters. The Bivariate Beta density of (ζ_c, ζ_T), has a complicated form, however, we can work with the density of the latent variables, (U, V, W):$p(\alpha, b|data, \Psi) \propto Unif(b|0,1)\Gamma(\alpha|a_\alpha, b_\alpha)$ $\prod_{l-1}^{L-1} Beta(U_l|\alpha, 1-b)Beta(V_l|\alpha, 1-b)Beta(W_l|1+\alpha-b, b)$. We sample from the proposal distribution, $(log(\alpha'), logit(b'))' \sim N_2((log(\alpha^{(b)}), logit(b^{(b)})), cS^2_{\alpha b})$, where $S^2_{\alpha b}$ is updated from the average variances and covariance of posterior samples of $((log(\alpha), logit(b))$ under initial runs, and c is updated from initial runs to optimize mixing. Draw $\eta \sim Unif(0,1)$.
If $\eta < min\left\{1, \frac{\Gamma(\alpha'|a_\alpha,b_\alpha)\prod_{l-1}^{L-1} Beta(U_l|\alpha',1-b')Beta(V_l|\alpha',1-b')Beta(W_l|1+\alpha'-b',b')\alpha' b'(1-b')}{\Gamma(\alpha^{(b)}|a_\alpha,b_\alpha)\prod_{l-1}^{L-1} Beta(U_l|\alpha^{(b)},1-b^{(b)})Beta(V_l|\alpha^{(b)},1-b^{(b)})Beta(W_l|1+\alpha^{(b)}-b^{(b)},b^{(b)})\alpha^{(b)}b^{(b)}(1-b^{(b)})}\right\}$
set $(\alpha^{(b+1)}, b^{(b+1)})' = (\alpha', b')'$, else $(\alpha^{(b+1)}, b^{(b+1)})' = (\alpha^{(b)}, b^{(b)})'$

APPENDIX B: Conditional Predictive Ordinate for gamma DDPMM

Here we provide the details of how we arrived to the expression necessary for computing the CPO values under the gamma DDPMM. As our data example in Section 4.1 does not contain any random covariates, we will derive the expression without covariates, however, the derivation can easily be extended to include random covariates in the curve-fitting setting.

Recall that the model, under the truncated version, can be written in hierarchical form as,

$$\begin{aligned}
t_{is}|\mathrm{w}_{is}, \boldsymbol{\theta} &\overset{ind}{\sim} \Gamma(t_{is}|\boldsymbol{\theta}_{\mathrm{w}_{is}}) \quad \text{for } i = 1, ..., n_s \quad s \in \{C, T\} \\
\mathbf{w}|\{\zeta_{lC}, \zeta_{lT}\} &\sim \prod_{s\in\{C,T\}} \prod_{i=1}^{n_s} \sum_{l=1}^{L} \left[(1-\zeta_{ls}) \prod_{r=1}^{l-1} \zeta_{rs}\right] \delta_l(\mathrm{w}_{is}) \\
\boldsymbol{\theta}_l|\boldsymbol{\mu}, \boldsymbol{\Sigma} &\overset{iid}{\sim} N_2(\boldsymbol{\theta}_l|\boldsymbol{\mu}, \boldsymbol{\Sigma}) \\
(\zeta_{lC}, \zeta_{lT})|\alpha, b &\overset{iid}{\sim} Biv - Beta((\zeta_{lC}, \zeta_{lT})|\alpha, b) \quad \text{for} l = 1, ..., L-1
\end{aligned}$$

with priors, $\alpha \sim \Gamma(\alpha|a_\alpha, b_\alpha)$, $b \sim Unif(b|0,1)$, $\boldsymbol{\mu} \sim N_2(\boldsymbol{\mu}|a_\mu, B_\mu)$, and $\boldsymbol{\Sigma} \sim IWish(\boldsymbol{\Sigma}|a_\Sigma, B_\Sigma)$. Let $\Psi = (\alpha, b, \boldsymbol{\mu}, \boldsymbol{\Sigma})$. The predictive density for a new survival time from group s, t_{0s}, is given by:

$$\begin{aligned}
p(t_{0s}|data) &= \int\int \Gamma(t_{0s}|\boldsymbol{\theta}_{\mathrm{w}_{0s}}) \left(\sum_{l=1}^{L} p_{ls}\delta_l(\mathrm{w}_{0s})\right) p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}, \Psi|data) d\mathrm{w}_{0s} d\boldsymbol{\theta} d\mathbf{w} d\boldsymbol{p} d\Psi \\
&= \int \left(\sum_{l=1}^{L} p_{ls}\Gamma(t_{0s}|\boldsymbol{\theta}_l)\right) p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}, \Psi|data) d\boldsymbol{\theta} d\mathbf{w} d\boldsymbol{p} d\Psi
\end{aligned}$$

Let s' be the experimental group that s is not, $data = \{\boldsymbol{t}_s, \boldsymbol{t}_{s'}\}$, and A be the normalizing constant for $p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}, \Psi | data)$. Namely, $p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}, \Psi | data) = \left[\{\prod_{i=1}^{n_s} \Gamma(t_{is}|\boldsymbol{\theta}_{\mathrm{w}_{is}})\} \left\{\prod_{i=1}^{n_{s'}} \Gamma(t_{is'}|\boldsymbol{\theta}_{\mathrm{w}_{is'}})\right\} p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}, \Psi)\right] / \left[\int \{\prod_{i=1}^{n_s} \Gamma(t_{is}|\boldsymbol{\theta}_{\mathrm{w}_{is}})\} \left\{\prod_{i=1}^{n_{s'}} \Gamma(t_{is'}|\boldsymbol{\theta}_{\mathrm{w}_{is'}})\right\} p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}, \Psi) d\boldsymbol{\theta} d\mathbf{w} d\boldsymbol{p} d\Psi\right]$. Note that $p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}, \Psi) = N_2(\boldsymbol{\theta}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) \left(\prod_{i=1}^{n_s} \sum_{l=1}^{L} p_{ls} \delta_l(\mathrm{w}_{is})\right) \left(\prod_{i=1}^{n_{s'}} \sum_{l=1}^{L} p_{ls'} \delta_l(\mathrm{w}_{is'})\right) Biv-Beta(\boldsymbol{p} \equiv (\zeta_s, \zeta_{s'})|\alpha, b) \Gamma(\alpha|a_\alpha, b_\alpha) Unif(b|0,1) N_2(\boldsymbol{\mu}|a_\mu, B_\mu) IWish(\boldsymbol{\Sigma}|a_\Sigma, B_\Sigma)$.

The CPO of the ith survival time in group s is defined as, $CPO_{is} = p(t_{is}|\boldsymbol{t}_{(-i)s}, \boldsymbol{t}_{s'}) = \int \Gamma(t_{is}|\boldsymbol{\theta}_{\mathrm{w}_{0s}}) \left(\sum_{l=1}^{L} p_{ls} \delta_l(\mathrm{w}_{0s})\right) p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}_{(-i)s}, \Psi) d\boldsymbol{\theta} d\mathbf{w}_{(-i)s} d\boldsymbol{p} d\Psi d\mathrm{w}_{0s}$, where $\mathbf{w}_{(-i)s}$ is the vector $\mathbf{w}$ with the i^{th} member of group s removed. Similarly, $data_{(-i)s}$ represents $data$ with the i^{th} member in group s removed. Now, consider $p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}_{(-i)s}, \Psi | data_{(-i)s}) =$

$$\frac{p(data_{(-i)s}|\boldsymbol{\theta}, \mathbf{w}_{(-i)s}) p(\boldsymbol{\theta}, \mathbf{w}_{(-i)s}, \boldsymbol{p}, \Psi)}{\int p(data_{(-i)s}|\boldsymbol{\theta}, \mathbf{w}_{(-i)s}) p(\boldsymbol{\theta}, \mathbf{w}_{(-i)s}, \boldsymbol{p}, \Psi) d\mathbf{w}_{(-i)s} d\boldsymbol{p} d\Psi}$$
$$= \frac{\left\{\prod_{j \neq i}^{n_s} \Gamma(t_{js}|\boldsymbol{\theta}_{\mathrm{w}_{js}})\right\} \left\{\prod_{i=1}^{n_{s'}} \Gamma(t_{is'}|\boldsymbol{\theta}_{\mathrm{w}_{is'}})\right\} p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}_{(-i)s}, \Psi)}{\int \left\{\prod_{j \neq i}^{n_s} \Gamma(t_{js}|\boldsymbol{\theta}_{\mathrm{w}_{js}})\right\} \left\{\prod_{i=1}^{n_{s'}} \Gamma(t_{is'}|\boldsymbol{\theta}_{\mathrm{w}_{is'}})\right\} p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}_{(-i)s}, \Psi) d\boldsymbol{\theta} d\mathbf{w}_{(-i)s} d\boldsymbol{p} d\Psi}$$

Let B_{is} be the normalizing constant of $p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}_{(-i)s}, \Psi | data_{(-i)s})$:
$B_{is} = \int \left\{\prod_{j \neq i}^{n_s} \Gamma(t_{js}|\boldsymbol{\theta}_{\mathrm{w}_{js}})\right\} \left\{\prod_{i=1}^{n_{s'}} \Gamma(t_{is'}|\boldsymbol{\theta}_{\mathrm{w}_{is'}})\right\} p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}_{(-i)s}, \Psi) d\boldsymbol{\theta} d\mathbf{w}_{(-i)s} d\boldsymbol{p} d\Psi$
Then, we can write $p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}_{(-i)s}, \Psi | data_{(-i)s}) =$

$$\frac{\{\prod_{i=1}^{n_s} \Gamma(t_{is}|\boldsymbol{\theta}_{\mathrm{w}_{is}})\} \left\{\prod_{i=1}^{n_{s'}} \Gamma(t_{is'}|\boldsymbol{\theta}_{\mathrm{w}_{is'}})\right\} p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}, \Psi)}{B_{is} \Gamma(t_{is}|\boldsymbol{\theta}_{\mathrm{w}_{is}}) p(\mathrm{w}_{is}|\boldsymbol{p})} = \frac{A}{B_{is}} \frac{p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}, \Psi | data)}{\Gamma(t_{is}|\boldsymbol{\theta}_{\mathrm{w}_{is}}) p(\mathrm{w}_{is}|\boldsymbol{p})}$$

Thus, $CPO_{is} = \int \Gamma(t_{is}|\boldsymbol{\theta}_{\mathrm{w}_{0s}}) p(\mathrm{w}_{0s}|\boldsymbol{p}) p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}_{(-i)s}, \Psi) d\boldsymbol{\theta} d\mathbf{w}_{(-i)s} d\boldsymbol{p} d\Psi d\mathrm{w}_{0s} =$
$\int \Gamma(t_{is}|\boldsymbol{\theta}_{\mathrm{w}_{0s}}) \left(\int p(\mathrm{w}_{0s}, \mathrm{w}_{is}|\boldsymbol{p}) d\mathrm{w}_{is}\right) p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}_{(-i)s}, \Psi) d\boldsymbol{\theta} d\mathbf{w}_{(-i)s} d\boldsymbol{p} d\Psi d\mathrm{w}_{0s} = A/B_{is} \int [\Gamma(t_{is}|\boldsymbol{\theta}_{\mathrm{w}_{0s}})$
$p(\mathrm{w}_{0s}, \mathrm{w}_{is}|\boldsymbol{p})] / [\Gamma(t_{is}|\boldsymbol{\theta}_{\mathrm{w}_{is}}) p(\mathrm{w}_{is}|\boldsymbol{p})] \, p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}, \Psi | data) d\mathrm{w}_{0s} d\boldsymbol{\theta} d\mathbf{w} d\boldsymbol{p} d\Psi = A/B_{is} \int \left[\sum_{l=1}^{L} p_{ls} \Gamma(t_{is}|\right.$
$\left.\boldsymbol{\theta}_l)\right] / [\Gamma(t_{is}|\boldsymbol{\theta}_{\mathrm{w}_{is}})] \, p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}, \Psi | data) d\mathrm{w}_{0s} d\boldsymbol{\theta} d\mathbf{w} d\boldsymbol{p} d\Psi$. (Note: $p(\mathrm{w}_{0s}|\mathrm{w}_{is}, \boldsymbol{p}) = p(\mathrm{w}_{0s}|\boldsymbol{p})$). All that is left is to be able to evaluate A/B_{is}:

$$\begin{aligned}
\left(\frac{A}{B_{is}}\right)^{-1} &= \frac{1}{A} \int \left\{\prod_{j \neq i}^{n_s} \Gamma(t_{js}|\boldsymbol{\theta}_{\mathrm{w}_{js}})\right\} \left\{\prod_{i=1}^{n_{s'}} \Gamma(t_{is'}|\boldsymbol{\theta}_{\mathrm{w}_{is'}})\right\} \underbrace{\left(\int p(\mathrm{w}_{is}|\mathbf{w}_{(-i)s}, \boldsymbol{p}) d\mathrm{w}_{is}\right)}_{1} \\
&\quad \times p(\mathbf{w}_{(-i)s}|\boldsymbol{p}) p(\boldsymbol{p}, \boldsymbol{\theta}, \Psi) d\boldsymbol{\theta} d\mathbf{w}_{(-i)s} d\boldsymbol{p} d\Psi \\
&= \frac{1}{A} \int \left\{\prod_{j \neq i}^{n_s} \Gamma(t_{js}|\boldsymbol{\theta}_{\mathrm{w}_{js}})\right\} \left\{\prod_{i=1}^{n_{s'}} \Gamma(t_{is'}|\boldsymbol{\theta}_{\mathrm{w}_{is'}})\right\} p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}, \Psi) d\boldsymbol{\theta} d\mathbf{w} d\boldsymbol{p} d\Psi \\
&= \frac{1}{A} \int \frac{\left\{\prod_{j \neq i}^{n_s} \Gamma(t_{js}|\boldsymbol{\theta}_{\mathrm{w}_{js}})\right\} \left\{\prod_{i=1}^{n_{s'}} \Gamma(t_{is'}|\boldsymbol{\theta}_{\mathrm{w}_{is'}})\right\}}{\Gamma(t_{is}|\boldsymbol{\theta}_{\mathrm{w}_{is}})} p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}, \Psi) d\boldsymbol{\theta} d\mathbf{w} d\boldsymbol{p} d\Psi \\
&= \int \frac{1}{\Gamma(t_{is}|\boldsymbol{\theta}_{\mathrm{w}_{is}})} p(\boldsymbol{\theta}, \boldsymbol{p}, \mathbf{w}, \Psi) d\boldsymbol{\theta} d\mathbf{w} d\boldsymbol{p} d\Psi
\end{aligned}$$

References

Argiento, R., Guglielmi, A., and Pievatolo, A. (2009), "A comparison of nonparametric priors in hierarchical mixture modeling for AFT regression," *Journal of Statistical Planning and Inference*, 139, 3989–4005.

Chen, M.-H., Shao, Q., and Ibrahim, J. (2000), *Monte Carlo Methods in Bayesian Computation*, Statistics, Springer.

Conner, R. and Mosemann, J. (1969), "Concepts of independence for proportions with a generalization of the Dirichlet distribution," *Journal of the American Statistical Association*, 64, 194–206.

DeIorio, M., Müller, P., Rosner, G. L., and MacEachern, S. N. (2004), "An ANOVA Model for Dependent Random Measures," *Journal of the American Statistical Association*, 99, 205–215.

DeIorio, M., Müller, P., Rosner, W., and Rosner, G. L. (2009), "Bayesian Nonparametric Nonproportional Hazards Survival Modeling," *Biometrics*, 65, 762–771.

Ferguson, T. S. (1973), "A Bayesian analysis of some nonparametric problems," *The Annals of Statistics*, 1, 209–230.

Fronczyk, K. and Kottas, A. (2010), "A Bayesian Nonparametric Modeling Framework for Developmental Toxicity Studies," Tech. rep., University of California, Santa Cruz.

Geisser, S. and Eddy, W. (1979), "A Predictive Approach to Model Selection," *American Statistical Association*, 74, 153–160.

Gelfand, A., Kottas, A., and MacEachern, S. (2005), "Bayesian Nonparametric Spatial Modeling With Dirichlet Process Mixing," *Journal of the American Statistical Association*, 100, 1021–1035.

Gelfand, A. E. and Mallick, B. K. (1995), "Bayesian analysis of proportional hazards models built from monotone functions," *Biometrics*, 51, 843–852.

Hall, W. J. and Wellner, J. A. (1981), "Mean Residual Life," in *Statistics and Related Topics*, eds. Csörgö, M., Dawson, D., Rao, J., and Saleh, A., North-Holland Publishing Company.

Hanson, T. and Johnson, W. O. (2002), "Modeling regression error with a mixture of Pólya tress," *Journal of American Statistical Association*, 97, 1020–1033.

Hanson, T. E., Jara, A., and Zhao, L. (2012), "A Bayesian Semiparametric Temporally-Stratified Proportional Hazards Model with Spatial Frailties," *Bayesian Analysis*, 7, 147–188.

Ibrahim, J., Chen, M.-H., and Sinha, D. (2001), *Bayesian Survival Analysis*, Statistics, Springer.

Ibrahim, J. G., Chen, M.-H., and MacEachern, S. N. (1999), "Bayesian Variable Selection for Proportional Hazards Models," *The Canadian Journal of Statistics*, 27, 701–717.

Ishwaran, H. and James, L. F. (2001), "Gibbs Sampling Methods for Stick-Breaking Priors," *American Statistical Association*, 96, 161–173.

Kottas, A., Behseta, S., Moorman, D., Poynor, V., and Olson, C. (2012), "Bayesian Nonparametric Analysis of Neuronal Intensity Rates," *Journal of Neuroscience Methods*, 203, 241–253.

Kottas, A., DeYoreo, M., and Poynor, V. (2013), "Bayesian Nonparametric Inference - Why and How," *Bayesian Analysis*, contributed discussion on paper by P. Ḿuller and R. Mitra.

Kottas, A. and Gelfand, A. E. (2001), "Bayesian semiparametric median regression modeling," *Journal of American Statistical Association*, 96, 1458–1468.

Kottas, A. and Krnjajić, M. (2009), "Bayesian Semiparametric Modeling in Quantile Regression," *Scandinavian Journal of Statistics*, 36, 297–319.

MacEachern, S. N. (2000), "Dependent Dirichlet Processes," Tech. rep., Department of Statistics, Ohio State University.

McKenzie, E. (1985), "An autoregessive process for beta random variables," *Management Science*, 31, 988–997.

Michael, J. R. and Schucany, W. R. (2011), "The Mixture approach for simulating new families of bivariate Distributions with specified correlations," Tech. rep., Southern Methodist University.

Mudholkar, G. S. and Strivasta, D. K. (1993), "Exponentiated Weibull Family for Analyzing Bathtub Failure-Rate Data," *IEEE Transactions of Reliability*, 42, 299–302.

Müller, P., Erkanli, A., and West, M. (1996), "Bayesian Curve Fitting Using Multivariate Normal Mixtures," *Biometrika*, 83, 67–79.

Nadarajah, S. and Kotz, S. (2005), "Some bivariate beta distributions," *Journal of Theoretical and Applied Statistics*, 39, 457–466.

Olkin, I. and Liu, R. (2003), "A bivariate beta distribution," *Statistics and Probability Letters*, 62, 407–412.

Poynor, V. (2010), "Bayesian inference for mean residual life functions in survival analysis," Master's thesis, University of California, Santa Cruz.

Poynor, V. and Kottas, A. (2014), "Nonparametric Bayesian Inference for Mean Residual Life Functions in Survival Analysis," ArXiv:1411.7481 [stat.ME].

Rodriguez, A. and ter Horst, E. (2008), "Bayesian dynamic density estimation," *Bayesian Analysis*, 3, 339–366.

Sethuraman, J. (1994), "A Constructive Definition of Dirichlet Priors," *Statistica Sinica*, 4, 639–650.

Taddy, M. and Kottas, A. (2010), "A Bayesian nonparametric approach to inference for quantile regression," *Journal of Business and Economic Statistics*, 28, 357–369.

Walker, S. G. and Mallick, B. K. (1999), "A Bayesian semiparametric accelerated failure time model," *Biometrics*, 55, 477–483.

Ying, Z., Jung, S., and Wei, L. (1995), "Survival analysis with median regression models," *Journal of the American Statistical Association*, 90, 178–184.